

COGGRAPH: Building bridges between cognitive science and computer graphics

Kartik Chandra (kach@mit.edu)

MIT CSAIL, Cambridge, MA 02139

Katherine Collins (kmc61@cam.ac.uk)

University of Cambridge

Chris Kymn (cjkyrn@berkeley.edu)

University of California, Berkeley

Kushin Mukherjee (kmukherjee2@wisc.edu)

University of Wisconsin - Madison

Anne Harrington (annekh@mit.edu)

MIT CSAIL

Sean P. Anderson (seanpaul@stanford.edu)

Stanford University

Arnav Verma (arnavv@stanford.edu)

Stanford University

Judith Fan (jefan@stanford.edu)

Stanford University

Background and motivation

In recent years, the field of computer graphics has achieved its longstanding dream of photorealism: modern graphics algorithms produce images that are indistinguishable from reality. Much like art at the advent of photography, then, computer graphics is now turning its gaze to the beholder: researchers are increasingly looking to cognitive science to engineer new modes of visual expression. Recent work has sought to apply insights from cognitive science to a variety of traditional graphics topics: from taking a perceptual approach to perspective (Hertzmann, 2023), to studying the theory of mind behind animation (Chandra, Li, Tenenbaum, & Ragan-Kelley, 2023), to applying theories of abstraction learning to build tools for geometry processing (Jones, Guerrero, Mitra, & Ritchie, 2023).

At the same time, a wave of recent work in cognitive science has addressed fundamental questions about visual expression: for example, how humans understand and create sketches (Fan, Bainbridge, Chamberlain, & Wammes, 2023), shapes (Dehaene et al., 2022), and symbols (Hofer, Kirby, & Levy, 2023). The field has also benefited from tools and methods from computer graphics: differentiable rendering systems (Kulkarni, Kohli, Tenenbaum, & Mansinghka, 2015), game engine physics simulators (Battaglia, Hamrick, & Tenenbaum, 2013), and Monte Carlo methods (Chandra, Chen, et al., 2023) have become indispensable in modeling human perception and intuitive physics.

Recognizing this growing interdisciplinary exchange of ideas, we are proposing a workshop to begin building formal bridges between the cognitive science and computer graphics communities. By bringing these two disciplines into contact, we seek to address two related sets of questions:

Q1 (Graphics → Cognition): How can powerful tools, methods, and ideas from computer graphics help us study the human mind? How can new algorithms for differentiable rendering help us model vision as inverse graphics? How can the design of modern physics simulators help us better understand intuitive physics? How can the principles of animation and visual storytelling give insight into human motor control and social cognition? How can non-photorealistic rendering systems shed light on the human capacity for abstraction and pictorial representation?

Q2 (Cognition → Graphics): How can insight from cognitive science help us better articulate and approach the goals of computer graphics? How can the cognitive science behind diagrams, sketches, representation, and iconicity be put to work in designing new tools for visual communication—from diagrams in a medical textbook to data visualizations in a newspaper? How can the study of communication and abstraction be applied to invent new visual media for artists to express themselves in? How can insights about theory of mind and emotion inform the next generation of tools for storytellers and animators?

Approach and schedule

Because this workshop bridges two distant research communities, it will have a unique format. In the spring, we will convene a series of virtual pre-events featuring short talks and discussions between invited speakers from both cognitive science and computer graphics. Then, in the summer, we will host in-person activities at CogSci 2024.

Virtual pre-events Before the main conferences, we will host four virtual events, each organized around a pillar of cognitive science with deep connections to graphics: Perception, Reasoning, Communication, and Expression. Each event will feature speakers from both cognitive science and computer graphics. Speakers will be invited to give brief opening remarks on Q1 and Q2. Then, they will engage in a moderated panel discussion. The virtual events will be open to the public and recorded.

In-person activities We will begin the in-person session at CogSci by reviewing the virtual seminars via an edited video that summarizes key themes from the recorded seminars. Then, in-person participants will engage in small-group discussions facilitated by the organizers. Finally, participants will reconvene and report back with insights from their discussions.

We intend to build on the success of past CogSci workshops that have used a similar format to foster engaging and memorable discussions among participants (Mukherjee, Huey, Rogers, & Fan, 2022; Binder et al., 2023; Mattar, Fan, Vong, & Wong, 2023).

After the workshop We will synthesize themes from the virtual events and in-person discussions in a review article authored by the organizers along with the invited speakers and workshop participants. More broadly, we hope this workshop will begin to build formal bridges between the cognitive science and computer graphics research communities, sparking fresh research agendas and long-term interdisciplinary collaborations.

Invited speakers

We have confirmed the following speakers for the workshop (names are hyperlinked to homepages).

Session 1: Perception

- **Ilker Yildirim**, a **cognitive scientist**, is an Assistant Professor at Yale whose work seeks to model human perception as inverse graphics.
- **Wenzel Jakob**, a **graphics researcher**, is an Associate Professor at EPFL and a leading expert on differentiable rendering for physically-accurate inverse graphics.
- **Vivian Paulun**, a **cognitive scientist**, is a postdoctoral fellow at MIT whose uses physics simulation to study how humans perceive material properties of objects.
- **Zachary Ferguson**, a **graphics researcher**, is a postdoctoral associate at MIT who studies algorithms for fast, accurate physics simulation.
- **Tzu-Mao Li**, a **graphics researcher**, is an Assistant Professor at UC San Diego and a leading expert on efficient algorithms for differentiable visual computing.

Session 2: Reasoning

- **Tomer Ullman** (moderator), a **cognitive scientist**, is an Assistant Professor at Harvard who studies people’s intuitive theories of the world.
- **Daniel Ritchie**, a **graphics researcher**, is an assistant professor at Brown who studies 3D scene understanding with a focus on systems that aid in visual creativity.
- **Kelsey Allen**, a **cognitive scientist**, is a Senior Research Scientist at Google DeepMind who studies tool use in humans and machines.
- **Gilbert Bernstein**, a **graphics researcher**, is an Assistant Professor at the University of Washington, Seattle, who designs domain-specific languages for physical simulation and differentiable programming.
- **Felix Binder**, a **cognitive scientist**, is a PhD student at UC San Diego who studies how the visual structure of the environment guides planning.

Session 3: Communication

- **Gabriel Greenberg** (moderator), a **philosopher**, is an Associate Professor in the Department of Philosophy at UCLA studying iconic representation.
- **Maneesh Agrawala**, a **graphics researcher**, is a professor at Stanford University who studies how cognitive design principles can be used to improve the effectiveness of audio/visual media.

- **Cindy Xiong**, a **cognitive scientist**, is an Assistant Professor at Georgia Tech studying data visualization.
- **Zoya Bylinskii**, a **graphics researcher**, is a Senior Research Scientist at Adobe and an expert on visualizations, fonts, and infographics.

Session 4: Expression

- **Aaron Hertzmann**, a **graphics researcher**, is a Principal Scientist at Adobe who studies the design of novel, powerful tools for artists. He is also a painter.
- **Sarah Schwettman**, an **cognitive AI researcher**, is a Research Scientist at MIT. She is a practicing artist and teaches a course at MIT on vision in art and neuroscience.
- **Yael Vinker**, a **graphics researcher**, is a PhD student at Tel-Aviv University who focuses on generative models to enhance creativity. She is also a painter and sculptor.
- **Cassidy Curtis**, an **artist**, is a Senior Staff Visual Designer at Google Research who studies the cognitive basis of animation. Previously a supervising animator at DreamWorks, he animated characters in films like *Shrek*, *Madagascar*, and *How to Train Your Dragon*.

Organizers

This workshop is chaired by **Kartik Chandra**, a PhD student at MIT, **Anne Harrington**, a Master’s student at MIT, and **Judith Fan**, an Assistant Professor at Stanford.

The workshop is additionally co-organized by **Katherine Collins**, a PhD student at the University of Cambridge; Chris Kymn, a PhD student at UC Berkeley; **Kushin Mukherjee**, a PhD student at UW Madison; **Sean Anderson**, a PhD student at Stanford; and **Arnav Verma**, a researcher at Stanford.

Acknowledgments

We thank Josh Tenenbaum and Jonathan Ragan-Kelley for their advice and support in organizing this workshop.

References

- Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, 110(45), 18327–18332.
- Binder, F. J., Cross, L. M., Friedman, Y., Hawkins, R., Yamins, D. L., & Fan, J. E. (2023). Advancing cognitive science and ai with cognitive-ai benchmarking. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 45).
- Chandra, K., Chen, T., Li, T.-M., Ragan-Kelley, J., & Tenenbaum, J. (2023). Inferring the future by imagining the past. *NeurIPS*.
- Chandra, K., Li, T.-M., Tenenbaum, J., & Ragan-Kelley, J. (2023). Acting as inverse inverse planning. In *Acm siggraph 2023 conference proceedings* (pp. 1–12).
- Dehaene, S., Al Roumi, F., Lakretz, Y., Planton, S., & Sablé-Meyer, M. (2022). Symbols and mental programs: a hypothesis about human singularity. *Trends in Cognitive Sciences*.

- 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.
- Hertzmann, A. (2023). Toward a theory of perspective perception in pictures. *psyarxiv.com/c7vdr*.
- Hofer, M., Kirby, S., & Levy, R. (2023). Simplicity and informativeness in the evolution of combinatorial structure. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 45).
- Jones, R. K., Guerrero, P., Mitra, N. J., & Ritchie, D. (2023). Shapecoder: Discovering abstractions for visual programs from unstructured primitives. *arXiv preprint arXiv:2305.05661*.
- Kulkarni, T. D., Kohli, P., Tenenbaum, J. B., & Mansinghka, V. (2015). Picture: A probabilistic programming language for scene perception. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 4390–4399).
- Mattar, M. G., Fan, J. E., Vong, W. K., & Wong, L. (2023). How does the mind discover useful abstractions? In *Proceedings of the annual meeting of the cognitive science society* (Vol. 45).
- Mukherjee, K., Huey, H., Rogers, T. T., & Fan, J. E. (2022). From images to symbols: Drawing as a window into the mind. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 44).