

Information Theory and Cognitive Science

Thomas A. Langlois
MIT and New York University

Nathaniel Imel
New York University

Clara Meister
ETH Zurich

Eleonora Gualdoni
Apple Machine Learning Group

Daniel Polani
University of Hertfordshire

Noga Zaslavsky
New York University

Workshop Website: <https://sites.google.com/view/infocog-cogsci-2025/>

Keywords: Information Theory, Computational Cognitive Science, Artificial Intelligence

Overview and motivation

In his paper on “the cognitive revolution,” George Miller dates the “moment of conception of cognitive science” to a 1956 symposium organized by the ‘Special Interest Group in Information Theory’ (Miller, 2003, pp. 142). This historical moment reflects the shared intellectual roots of information theory, which was founded by Claude Shannon (1948), and cognitive science. However, the initial excitement about applications of information theory to cognition and many other fields was quickly turned off by critics, including Shannon himself, who argued that the theory’s scope outside the realm of engineering is limited (Shannon, 1956; Luce, 2003).

Yet, several decades later, we now see a renewed surge of interest in information theory, with numerous successful applications, ranging from theoretical neuroscience (e.g., Friston, 2010; Palmer et al., 2015; Tkačik & Bialek, 2016) to working memory (e.g., Bates et al., 2019; Jakob & Gershman, 2023), perception (e.g., Sims, 2016, 2018), decision making (e.g., Tishby & Polani, 2011; Lai & Gershman, 2021; Bhui et al., 2021), behavioral economics (e.g., Caplin et al., 2022; Prat-Carrabin & Woodford, 2022; Azeredo da Silveira et al., 2024), and language (e.g., Zaslavsky et al., 2018; Gibson et al., 2019), as well as broad applications in artificial intelligence (e.g., Tishby & Zaslavsky, 2015; Alemi et al., 2017; Du et al., 2020; Gualdoni et al., 2024). This recent body of work suggests that, in contrast to the initial reservations by Shannon and others, information theory provides a powerful framework for understanding both human and artificial intelligence at multiple levels of analysis (Marr, 1982). However, so far, information-theoretic approaches have been applied mostly independently in each area of cognitive science, without much interaction across disciplines.

The goal of this workshop is to create a multidisciplinary space for discussing the most recent advances at the intersection of information theory and cognitive science and to explore how this emerging research area can help the field advance toward a more comprehensive and principled mathematical theory of human cognition. In particular, we aim to attract and inspire intensive and meaningful dialogue

between researchers who, while active in a gamut of fields, share the vision of communicating across the domain via concepts and ideas from information theory.

Approach and workshop structure

To achieve our goals, we confirmed a diverse set of 4 invited talks by notable researchers in the field, representing a wide range of view points, from perception, cognitive neuroscience and decision making, to child development, and language. See Table 1 below for the full list of confirmed speakers, their planned topic for the talk, and our proposed tentative schedule. Each talk will be 30 minutes (including 5-10 minutes for questions) and will illuminate how information theory can help advance our understanding of one or more key aspects of human cognition. The talks will be followed by a panel discussion, featuring all invited speakers, to integrate the different perspectives and themes explored here. In addition to providing a broad overview of active research across different subdisciplines of cognitive science, we will discuss topics related to the current gaps between information-theoretic hypotheses, the efficiency of biological systems, and modern artificial intelligence. We crystallize the motto for the panel into the following question: “*Who needs compressive brains if you have expansive LLMs?*” We intend to encourage audience participation during the panel and believe this motto will not only facilitate engagement but also drive the discussion toward postulating promising avenues for future work. Altogether, the workshop will offer a forum to discuss cutting-edge applications of information theory to cognitive science and AI.

Target audience

We anticipate that this workshop would be of interest to a broad range of researchers in the cognitive science community, including, but limited to, those who study language, decision-making, perception, learning, and AI. Our invited speakers and their research activities reflect this range, with information-theoretic principles at the core of their current work. While the topic of the workshop would naturally appeal to those who are interested in computational models, we aim to make the workshop accessible and welcoming to anyone who is interested in understanding cognition. This workshop would be part of a series of related workshops on information theory and cognition. The first two

Time	Presenter	Topic
8:30	Nori Jacoby	<i>Efficient coding in auditory and visual representations</i>
9:00	Alison Gopnik	<i>Empowerment and causal learning during development</i>
9:30	Mimi Liljeholm	<i>Information theory, learning and decision making</i>
10:00		Break
10:30	Terry Regier	<i>Efficient communication, cross-linguistic semantic variation</i>
11:00	Panel	<i>Who needs compressive brains if you have expansive LLMs?</i>

Table 1: Tentative schedule for the workshop, see sites.google.com/view/infocog-cogsci-2025/schedule for final schedule.

workshops took place as part of the Annual Conference on Neural Information Processing Systems (NeurIPS) in 2022 and 2023, focusing primarily on information theory as a bridge between human and artificial cognition. By bringing this workshop to CogSci 2025, we hope to extend its scope and breadth across the cognitive science community.

Invited speakers and panelists

Nori Jacoby is an Assistant Professor of Psychology at Cornell University. He studies the internal representations that support and shape our sensory and cognitive abilities, combining techniques from machine learning and information theory with behavioral experiments and fieldwork in various locations around the globe. His talk will focus on mapping the geometry of internal representations using a new experimental paradigm involving sampling with people.

Alison Gopnik is a Professor of Psychology at the UC Berkeley. She is an expert in cognitive development, causal learning, and Bayesian models. In her talk, she will discuss recent work on how the concept of empowerment, an information-theoretic intrinsic reward signal, could bridge classical Bayesian causal learning and reinforcement learning, and may help to characterize causal learning in humans and enable it in machines.

Mimi Liljeholm is an Associate Professor of Cognitive Sciences at UC Irvine. Her research addresses how intelligent systems construct and deploy mental models of the world, with a focus on the roles of compositional generalization and intrinsic motivation. In her talk, she will demonstrate the use of Bayesian and Information Theoretic variables to describe core aspects of learning and decision-making, at behavioral and neural levels. She will then discuss alternative

approaches and reflect on some advantages and perils of computational cognitive modeling.

Terry Regier is a Professor of Linguistics and Cognitive Science at UC Berkeley. His research investigates the relation of language and cognition, through computational principles, cross-language semantic data, and behavioral experiments. In his talk, he will discuss information-theoretic approaches to cross-language variation in semantic categories (e.g., word meanings), focusing on communicative needs and implications to the relation of language and culture.

Organizers

Noga Zaslavsky is an Assistant Professor of Psychology at NYU. Her research aims to understand language and cognition from first principles, building on information theory and machine learning. She organized the two earlier workshops on information theory and cognition at NeurIPS 2022 and 2023.

Thomas A. Langlois is a Postdoctoral Associate in MIT’s Brain and Cognitive Sciences Department, jointly affiliated with NYU’s Department of Psychology, where he is currently working on applications of information theory to visual perception and language.

Nathaniel Imel is a PhD student in Cognition and Perception at NYU, where he is working on combining information theory and evolutionary dynamics in studying the emergence of efficient concepts and categories.

Clara Meister is a Postdoctoral Associate at ETH Zürich, developing material for a continuing studies program on AI and Data Technologies. Her research foci include language generation, psycholinguistics, and application of statistical methods and information theory to natural language processing.

Eleonora Gualdoni is an AIML resident at Apple, Machine Learning Research, in the Paris team led by Marco Cuturi.¹ She studies how humans connect language to visual inputs, as well as how deep-learning multimodal model perform on linguistic and visual tasks, aiming to enhance their human-like capabilities.

Daniel Polani is a Professor of Artificial Intelligence at the University of Hertfordshire, UK. He studies fundamental principles of flexible decision-making in animals and humans in complex environments through the informational lens and how to transfer these to artificial and robotic systems.

¹Workshop proposal finished prior to joining Apple

References

- Alemi, A. A., Fischer, I., Dillon, J. V., & Murphy, K. (2017). Deep variational information bottleneck. In *International conference on learning representations*.
- Azeredo da Silveira, R., Sung, Y., & Woodford, M. (2024, October). Optimally imprecise memory and biased forecasts. *American Economic Review*, *114*(10), 3075–3118. doi: 10.1257/aer.20201806
- Bates, C. J., Lerch, R. A., Sims, C. R., & Jacobs, R. A. (2019). Adaptive allocation of human visual working memory capacity during statistical and categorical learning. *Journal of vision*, *19*(2), 11–11.
- Bhui, R., Lai, L., & Gershman, S. J. (2021). Resource-rational decision making. *Current Opinion in Behavioral Sciences*, *41*, 15–21.
- Caplin, A., Dean, M., & Leahy, J. (2022). Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political Economy*, *130*(6), 1676–1715. doi: 10.1086/719276
- Du, Y., Tiomkin, S., Kiciman, E., Polani, D., Abbeel, P., & Dragan, A. (2020). AvE: Assistance via empowerment. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, & H. Lin (Eds.), *Advances in neural information processing systems* (Vol. 33, pp. 4560–4571). Curran Associates, Inc.
- Friston, K. (2010, Feb 01). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, *11*(2), 127–138.
- Gibson, E., Futrell, R., Piantadosi, S. P., Dautriche, I., Mahowald, K., Bergen, L., & Levy, R. (2019). How efficiency shapes human language. *Trends in Cognitive Sciences*, *23*(5), 389–407. doi: <https://doi.org/10.1016/j.tics.2019.02.003>
- Gualdoni, E., Tucker, M., Levy, R. P., & Zaslavsky, N. (2024). Bridging semantics and pragmatics in information-theoretic emergent communication. In *The thirty-eighth annual conference on neural information processing systems*.
- Jakob, A. M., & Gershman, S. J. (2023, jul). Rate-distortion theory of neural coding and its implications for working memory. *eLife*, *12*, e79450. doi: 10.7554/eLife.79450
- Lai, L., & Gershman, S. J. (2021). Policy compression: An information bottleneck in action selection. In *Psychology of learning and motivation* (Vol. 74, pp. 195–232). Elsevier.
- Luce, R. D. (2003). Whatever happened to information theory in psychology? *Review of General Psychology*, *7*(2), 183–188. doi: 10.1037/1089-2680.7.2.183
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York, NY, USA: Henry Holt and Co., Inc.
- Miller, G. A. (2003). The cognitive revolution: a historical perspective. *Trends in Cognitive Sciences*, *7*(3), 141–144.
- Palmer, S. E., Marre, O., Berry, M. J., & Bialek, W. (2015). Predictive information in a sensory population. *Proceedings of the National Academy of Sciences*, *112*(22), 6908–6913.
- Prat-Carrabin, A., & Woodford, M. (2022). Efficient coding of numbers explains decision bias and noise. *Nature Human Behaviour*, *6*(8), 1142–1152.
- Shannon, C. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, *27*(3), 379–423.
- Shannon, C. (1956). The bandwagon. *IRE Transactions on Information Theory*, *2*(1), 3–3. doi: 10.1109/TIT.1956.1056774
- Sims, C. R. (2016). Rate–distortion theory and human perception. *Cognition*, *152*, 181 - 198. doi: <http://dx.doi.org/10.1016/j.cognition.2016.03.020>
- Sims, C. R. (2018). Efficient coding explains the universal law of generalization in human perception. *Science*, *360*(6389), 652–656.
- Tishby, N., & Polani, D. (2011). Information theory of decisions and actions. In V. Cutsuridis, A. Hussain, & J. G. Taylor (Eds.), *Perception-action cycle: Models, architectures, and hardware* (pp. 601–636). New York, NY: Springer New York.
- Tishby, N., & Zaslavsky, N. (2015). Deep learning and the information bottleneck principle. In *2015 IEEE information theory workshop (ITW)* (pp. 1–5).
- Tkačik, G., & Bialek, W. (2016). Information processing in living systems. *Annual Review of Condensed Matter Physics*, *7*(1), 89–117. doi: 10.1146/annurev-conmatphys-031214-014803
- Zaslavsky, N., Kemp, C., Regier, T., & Tishby, N. (2018). Efficient compression in color naming and its evolution. *Proceedings of the National Academy of Sciences*, *115*(31), 7937–7942.