

Compositionality in minds, brains and machines: a unifying goal that cuts across cognitive sciences

Barbara Pomiechowska¹; Rachel Dudley²; Lio Wong³; Mathias Sablé-Meyer⁴

¹University of Birmingham; ²Central European University; ³MIT; ⁴Sainsbury Wellcome Centre, UCL

Keywords: Compositionality; Abstraction; Generalization;
Cognition
<https://workshop-compositionality-cogsci-2024.info/>

Overview and Motivation

Compositionality, or the ability to build complex representations from discrete elements, is an essential ingredient of human intelligence. Compositionality enables people to think productively, learn fast from limited experience, and generalize knowledge to new contexts without re-learning from scratch. It is also essential in information processing systems to efficiently represent structured data and has seen application in compression and symbolic Artificial Intelligence (AI). Historically, the notion of compositionality played a central role in linguistics and philosophy of mind. More recently, it is attracting a surge of interest throughout the domains of cognitive science. Compositional processes are leveraged for elucidating the nature of mental representations in cognition (Dehaene et al., 2022), understanding the functional organisation of the brain (Agrawal et al., 2019), or building AI systems that are robust to changes in the environment (Hupkes et al., 2020).

In experimentally addressing the mechanisms of compositionality in the mind, some researchers focused on natural language (Pylkkänen, 2019) and its precursors in animal communication (Zuberbühler, 2020; Schlenker et al., 2014). Others prioritized non-linguistic structures such as geometric shapes (Sablé-Meyer et al., 2021), functions (Schulz et al., 2017), music (Koelsch et al., 2002). These different empirical foci reflect two distinct ways of thinking about compositional thought: as a product of a generative system for expressing meaning (Szabó, 2012; Chomsky, 2006), versus a product of generative system to model mental representations (Fodor, 1975; Dehaene et al., 2022; Lake et al., 2015), and make apparent that a unified understanding of the origins and nature of compositional processes *is* yet to emerge. The increased ability to simultaneously record hundreds of single neurons across different brain regions has empowered systems neuroscience to target behaviors more complex than ever before, which holds the promise to adjudicate between theories that make indistinguishable computational predictions.

Recent empirical evidence suggests that certain forms of compositional thought emerge very early in life, and can be observed both in linguistic and non-linguistic contexts. For example, human infants spontaneously deploy compositional inferences to interpret speech (de Carvalho et al., 2021; Pomiechowska et al., 2022) and keep track of hidden objects in the absence of verbal prompts (Cesana-Arlotti et al., 2018). In contrast, despite recent spectacular advances in AI, compositionality does not seem to automatically emerge from very

large data sets (Marcus et al., 2022). Most artificial systems do not seem to use compositionality as resourcefully and flexibly as humans, as evidenced by their inability to generalize compositional models across domains and tasks (Bowers et al., 2023). What accounts for these striking differences between human infants and machines? What cognitive tools are available to young children and absent in artificial systems?

The goal of this workshop is to explore the possibility of a common conceptual framework capable of unifying research on compositionality across disciplinary and methodological boundaries. To achieve this goal, the workshop will be a forum for discussion with leading researchers in developmental and cognitive psychology, human and systems neuroscience, and AI. The speakers will share their latest research and engage with the following inter-related outstanding questions:

- How to evaluate whether a system is compositional? What is the smallest interesting test case for compositionality?
- What are each discipline's most recent empirical findings that can help build a shared understanding of compositional systems?
- What neural mechanisms underlie compositional cognitive processes?
- Is compositional thought implemented via one domain-general system or multiple domain-specific systems? What are its representational and architectural prerequisites?
- Is there a risk that a generous definition of compositionality makes it too universal to be useful?

Approach and Workshop Structure

Compositionality is both a foundational principle in linguistics and philosophy as well as an exciting re-emerging object of interest within several of cognitive science fields, in particular, psychology, neuroscience, and AI. Our aim is to provide the broad CogSci community with the opportunity for an in-depth interdisciplinary discussion on the nature of compositional processes, grounded in the latest empirical advances. We hope that this workshop will inspire novel integrative work across disciplines.

To maximize the scope of the workshop, it will have two components: an in-person full day event and a virtual seminar. The **in-person program** will be held at CogSci and will span a full day, comprising four sessions of talks and a panel discussion. The first session *Frameworks for Compositionality* will introduce overarching theoretical approaches. The subsequent sessions (*Compositionality in Cognition*; *Compositionality in Brains and Neurons*, *Compositionality in AI*) will focus on the latest empirical evidence in studying compositionality across cognitive sciences. We will conclude with

a panel discussion that will tie together various perspectives and findings, and explore future directions for research on compositional systems. We will encourage the panelists to engage with the key questions outlined above. The **virtual seminar**, *Compositionality in Mathematics and Nonlinguistic Reasoning*, will take place before the conference and will be dedicated to survey an emerging line of research on compositional thought beyond natural language, in humans, non-human animals, and computational models. Speakers will also join the in-person panel discussion.

Organizers and Speakers

Organizers will not be giving full presentations.

Rachel Dudley (organizer) is a Postdoctoral Fellow at the Central European University. Her work explores the precursors of semantic development in infancy.

Barbara Pomiechowska (organizer) is an Assistant Professor of Psychology at the University of Birmingham. She studies how symbolic representation and compositional thought emerge in human infants and young children.

Mathias Sablé-Meyer (organizer) is a Senior Research Fellow in Cognitive Neuroscience at the Sainsbury Wellcome Center, UCL. He works on the representation of abstract concepts in humans by studying geometric shapes and structured sequences, bridging the gap between computational description and neural implementations.

Lio Wong (organizer) is a PhD Candidate in Brain and Cognitive Sciences at MIT. Their research focuses on computational and cognitive models that integrate structured conceptual reasoning with language, and that learn new concepts and abstractions from language.

Fosca Al Roumi is a Researcher at NeuroSpin. She studies how the human brain represents structured sequences.

Lorenzo Ciccione is a Postdoctoral Researcher at NeuroSpin. He studies the perception of plots in humans.

Isabelle Dautriche is a Research Scientist at the French National Centre for Scientific Research. She studies language development and its links to the structure of languages.

Mohamady El-Gaby is a Postdoctoral Neuroscientist in Oxford. He investigates, at the cellular level, how abstract cognitive maps form and how they interact with spatial maps to guide flexible behaviour.

Kevin Ellis is an Assistant Professor of Computer Science at Cornell University. He builds AI systems that represent knowledge in the form of symbolic code, and treat learning as program synthesis.

Noémi Éltető is a PhD student at the Max Planck Institute for Computational Neuroscience. She investigates how individuals extract patterns from noisy, high-order sequences that underlie motor skills or grammar.

Stephen Ferrigno is an Assistant Professor at the University of Wisconsin-Madison. He studies the evolutionary, developmental, and cultural origins of human thought.

Alla Karpova is a Senior Group Leader at Janelia. Her work aims to identify what neural computations allow animals to

In-person program

Frameworks for Compositionality

Brenden Lake	Compositional generalization through meta-learning neural networks
Liina Pykkänen	How do our brains order syntactic and semantic computations when no order is imposed from the input?

Compositionality in Cognition

Eric Schulz	TBA
Isabelle Dautriche	Non-linguistic compositionality in 9-month-old infants
Fosca Al Roumi	The human brain compresses information in memory using a Language of Thought

Compositionality in Brains and Neurons

Mohamady El-Gaby	A cellular basis for mapping behavioural structure
Alla Karpova	Towards the neural underpinnings of compositionality
Liping Wang	The control of sequence working memory in macaque frontal cortex.

Compositionality in AI

Noémi Éltető	Reusing action sequences for efficient planning
Kevin Ellis	Proposing experiments and acquiring concepts using language and code

On-line seminar

Compositionality in Mathematics and Nonlinguistic Reasoning

Akshita Srinivasan & Elizabeth Spelke	Young children's compositional reasoning about numbers
Stephen Ferrigno	How do children, adults, and non-human primates represent center-embedded sequences?
Lorenzo Ciccione & Théo Morfousse	Exploring mathematical composition in children through drawing and music
Paul Muhle-Karbe	Flexible representations for context-dependent navigation

dynamically adjust their behavior to the structure and demands of the environment.

Brenden Lake is an Assistant Professor of Psychology and Data Science at NYU. He studies machine and human intelligence.

Théo Morfousse is a PhD Student at NeuroSpin. He is interested in shape perception in humans.

Paul Muhle-Karbe is an Assistant Professor of computational cognitive neuroscience at the University of Birmingham. His work examines the building blocks of higher cognitive functions.

Liina Pykkänen is a Professor of Linguistics and Psychology at NYU. Her research uncovers how complex meaning emerges in the human brain.

Eric Schulz is a Research Group Leader at the Max Planck Institute for Biological Cybernetics. His group studies building blocks of intelligence.

Elizabeth Spelke is a Professor of Psychology at Harvard. Her work explores uniquely human cognitive and learning abilities.

Akshita Srinivasan is a PhD Candidate at Harvard. She studies how children learn compositional concepts.

Liping Wang is a Research Group Leader at the Chinese Academy of Sciences. He studies neural mechanisms underlying sequence learning and working memory.

References

- Agrawal, A., Hari, K., & Arun, S. (2019). Reading increases the compositionality of visual word representations. *Psychological Science, 30*(12), 1707–1723.
- Bowers, J. S., Malhotra, G., Dujmović, M., Montero, M. L., Tsvetkov, C., Biscione, V., ... others (2023). Deep problems with neural network models of human vision. *Behavioral and Brain Sciences, 46*, e385.
- Cesana-Arlotti, N., Martín, A., Téglás, E., Vorobyova, L., Cetnarski, R., & Bonatti, L. L. (2018). Precursors of logical reasoning in preverbal human infants. *Science, 359*(6381), 1263–1266.
- Chomsky, N. (2006). *Language and mind*. Cambridge University Press.
- de Carvalho, A., Crimon, C., Barrault, A., Trueswell, J., & Christophe, A. (2021). “look! it is not a bamoule!”: 18- and 24-month-olds can use negative sentences to constrain their interpretation of novel word meanings. *Developmental science, 24*(4), e13085.
- 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*.
- Fodor, J. A. (1975). *The language of thought* (Vol. 5). Harvard university press.
- Hupkes, D., Dankers, V., Mul, M., & Bruni, E. (2020). Compositionality decomposed: How do neural networks generalise? *Journal of Artificial Intelligence Research, 67*, 757–795.
- Koelsch, S., Gunter, T. C., Cramon, D. Y. v., Zysset, S., Lohmann, G., & Friederici, A. D. (2002). Bach speaks: a cortical “language-network” serves the processing of music. *Neuroimage, 17*(2), 956–966.
- Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science, 350*(6266), 1332–1338.
- Marcus, G., Davis, E., & Aaronson, S. (2022). A very preliminary analysis of dall-e 2. *arXiv preprint arXiv:2204.13807*.
- Pomiechowska, B., Brody, G., Téglás, E., & Kovács, Á. M. (2022). Combinatorial thought in infancy: Language processing reveals conceptual combination. *PsyArXiv* <https://doi.org/10.31234/osf.io/hekp5>.
- Pylkkänen, L. (2019). The neural basis of combinatory syntax and semantics. *Science, 366*(6461), 62–66.
- Sablé-Meyer, M., Fagot, J., Caparos, S., van Kerkoerle, T., Amalric, M., & Dehaene, S. (2021). Sensitivity to geometric shape regularity in humans and baboons: A putative signature of human singularity. *Proceedings of the National Academy of Sciences, 118*(16), e2023123118.
- Schlenker, P., Chemla, E., Arnold, K., Lemasson, A., Ouatara, K., Keenan, S., ... Zuberbühler, K. (2014). Monkey semantics: two ‘dialects’ of campbell’s monkey alarm calls. *Linguistics and Philosophy, 37*, 439–501.
- Schulz, E., Tenenbaum, J. B., Duvenaud, D., Speekenbrink, M., & Gershman, S. J. (2017). Compositional inductive biases in function learning. *Cognitive psychology, 99*, 44–79.
- Szabó, Z. G. (2012). The case for compositionality. In W. Hinzen, E. Machery, & M. Werning (Eds.), *The oxford handbook of compositionality*. Oxford: Oxford University Press.
- Zuberbühler, K. (2020). Syntax and compositionality in animal communication. *Philosophical Transactions of the Royal Society B, 375*(1789), 20190062.