

Symbolic and Sub-Symbolic Systems in People and Machines

Simon Valentin¹, Bonan Zhao², Chentian Jiang¹, Neil R. Bramley² & Christopher G. Lucas¹,

1. School of Informatics, University of Edinburgh; 2. Department of Psychology, University of Edinburgh

Keywords: symbolism; program induction; machine learning; mental representation

Overview

To what extent is symbolic processing required for intelligent behavior? Advances in sub-symbolic deep learning systems and explicitly symbolic probabilistic program induction approaches have recently reinvigorated this long standing question about cognition. This workshop will bring together established and newly-emerging perspectives on the debate and explore the recently rekindled interest in hybrid architectures.

The notion that intelligence rests on symbolic computations dominated early work in cognitive science (Newell & Simon, 1976). Symbolic approaches promised to explain the human capacity for reason, common sense, imagination, as well as linguistic competence. The broad idea was and is that symbol manipulation systems are uniquely capable of logic and reasoning, and enable compositionality, allowing thought to exhibit the systematicity and productivity that enable natural languages to make “infinte use of finite means” (von Humboldt, 1836). However, purely symbolic systems are yet to live up to this promise, facing foundational issues dealing with uncertainty and scalability, as well as unresolved questions around the origin and nature of the symbols, how they can be imputed from continuous sensory inputs, not to mention how such systems can be implemented in a distributed neural system. Indeed much of the maturation of cognitive science over the last 40 years has involved a move away from “logicism” toward probabilistic inference (Oaksford, Chater, et al., 2007), and neurally plausible reinforcement-driven computations (Schultz, Dayan, & Montague, 1997).

Sub-symbolic perspectives have recently received considerable attention due to the successes of deep learning in applied domains and have a long history as a proposed alternative to symbolic systems (LeCun, Bengio, & Hinton, 2015). However, their shortcomings also point to fundamental challenges associated with a purely sub-symbolic approach (Lake, Ullman, Tenenbaum, & Gershman, 2017). For instance, deep learning systems can be confused by subtle changes to the data distribution and may fail to enforce trivial forms of compositionality in their generalizations.

Recently, a third path has regained traction, with a glut of hybrid architectures being proposed that seek to combine symbolic and sub-symbolic approaches (Nye, Solar-Lezama,

Tenenbaum, & Lake, 2020; Valkov, Chaudhari, Srivastava, Sutton, & Chaudhuri, 2018). These may (at least superficially) reflect common psychological theory distinctions between low level processing —associated with perception and rapid, intuitive responses—and thinking—slower, effortful conscious processing. However, these connections are yet to be fully unpacked, and many challenging technical problems remain for these approaches.

Goals and Scope

Our workshop will explore:

- Connections between advances in neuro-symbolic machine learning systems and theories of cognition.
- Potential for hybrid systems to exhibit intelligent behavior.
- Limits of solely symbolic or sub-symbolic systems.
- Which forms of hybrid system, if any, might best characterize human cognition.
- How to design experiments to distinguish between symbolic and sub-symbolic processing in cognition.

Target Audience

We expect this workshop to be of broad interest to most of the CogSci audience, including cognitive and developmental psychologists, neuroscientists, linguists, machine learning researchers, roboticists and philosophers.

Organizers and Presenters

The following organizers, speakers and panelists have confirmed their attendance.

Simon Valentin (Organizer) is a PhD student with Chris Lucas and Neil Bramley. His research focuses on causal inference and active learning.

Bonan Zhao (Organizer) is a PhD student with Neil Bramley. She studies the interplay between causality and generalization using computational models.

Chentian Jiang (Organizer) is a PhD student with Chris Lucas. Her research focuses on topics in causal learning, transfer learning, and program induction

Neil R. Bramley (Organizer) is Lecturer in Cognitive Psychology at the University of Edinburgh with a focus on causal cognition, social cognition, active learning and discovery.

Christopher G. Lucas (Organizer) is Lecturer in Informatics at the University of Edinburgh, with a focus on computational models of causal, inductive and active learning.

Peter Dayan is a Rumelhart Prize winning theoretical neuroscientist and Director of the Max Planck Institute for Biological Cybernetics in Tübingen. His research focuses on the computational and neural foundations of learning and decision making.

Nick Chater is Professor of Behavioural Science at Warwick Business School. His research focuses on experimental, computational and mathematical analyses of cognitive processes, especially reasoning, decision making and language.

Jessica Hamrick is a Senior Research Scientist at DeepMind. Her work combines cognitive science, model-based deep reinforcement learning, and planning to build machines that can flexibly build and deploy models of the world.

Brenden Lake is Assistant Professor at NYU and Research Scientist at Facebook AI Research. He builds computational models of everyday cognitive abilities, focusing on problems that are easier for people than they are for machines.

Eric Schulz is a Research Group Leader at the Max Planck Institute for Biological Cybernetics in Tübingen. His lab builds computational models of human intelligence to study how people learn, generalize and explore.

Kevin Ellis is Assistant Professor at Cornell University. His research involves building machine learning systems that generalize effectively.

Steven Piantadosi is Assistant Professor at UC Berkeley. His research focuses on how people learn language and create conceptual systems

Jiajun Wu is Assistant Professor at Stanford University. His research draws inspiration from human cognition to study machine perception and reasoning about the physical world.

Judith Fan is Assistant Professor at UC San Diego. Her lab builds computational models of perception, memory, motor planning and social cognition.

Thomas Icard is Assistant Professor at Stanford University. He develops formal models of human cognitive abilities such as language, explanation, and causal reasoning.

Guillermo Puebla is a postdoc at the University of Bristol. He studies relational reasoning and analogical generalization through simulations with artificial neural networks, analogical inference models, and hybrid symbolic-neural systems.

Kelsey Allen is a PhD student with Josh Tenenbaum at MIT. Her work is focused on the interactions between predictive representations and planning in human cognition.

Kate McCurdy is a PhD student at the University of Edinburgh. Her research focuses on computational models of natural language use.

Workshop Structure

This workshop is scheduled to span a full day. All talks will last 20 minutes, including time for questions from the audience and for the change over between speakers. The final part will be a panel discussion led by the organizers,

where the panelists will tie together and contrast themes encountered throughout the day in conversation with the audience. For the final schedule, see the workshop website at <https://sassy-2021.github.io>.

Presenter	Topic
Valentin & Lucas	Introductory remarks
Brenden Lake	Probabilistic program induction
Jessica Hamrick	Model-based deep reinforcement learning
Kevin Ellis	Neurosymbolic learning and program synthesis
Eric Schulz	Multi-task reinforcement learning
Kelsey Allen	Symbolic planning and simulation in the physical world
Steven Piantadosi	Some ideas in program induction
Jiajun Wu	The neurosymbolic concept learner
Judith Fan	Deep neural representations for visual communication
Neil Bramley	Program synthesis in development
Guillermo Puebla	Relational reinforcement learning
Kate McCurdy	Neural encoder-decoder networks as cognitive models of grammar
Bonan Zhao	Object-based causal generalization
Thomas Icard	Program induction and causal inference
Peter Dayan	Panel discussion
Nick Chater	Panel discussion

Acknowledgments

Bramley and Lucas are part-supported by an UKRI EPSRC New Investigator Grant on “computational constructivism” (EP/T033967/1)

References

- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- Newell, A., & Simon, H. A. (1976, March). Computer science as empirical inquiry: Symbols and search. *Commun. ACM*, 19(3), 113–126.
- Nye, M. I., Solar-Lezama, A., Tenenbaum, J. B., & Lake, B. M. (2020). Learning compositional rules via neural program synthesis. *arXiv preprint arXiv:2003.05562*.
- Oaksford, M., Chater, N., et al. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford University Press.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593–1599.
- Valkov, L., Chaudhari, D., Srivastava, A., Sutton, C., & Chaudhuri, S. (2018). Houdini: Lifelong learning as program synthesis. In *Advances in Neural Information Processing Systems* (pp. 8687–8698).
- von Humboldt, W. (1836). *On Language*. Cambridge University Press.