UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

How Does Current AI Stack Up Against Human Intelligence?

Permalink

Journal

Authors

Forbus, Ken
Gentner, Dedre
Laird, John E.
et al.

Publication Date

2019

Peer reviewed

How Does Current AI Stack Up Against Human Intelligence?

Ken Forbus (forbus@northwestern.edu) and Dedre Gentner (gentner@northwestern.edu)
Department of Electrical Engineering and Computer Science, & Department of Psychology, Northwestern University
Evanston, IL. 60208 USA

John E. Laird (laird@umich.edu)
Department of Electrical Engineering and Computer Science, University of Michigan
Ann Arbor, MI. 48109 USA

Thomas Shultz (thomas.shultz@mcgill.ca) and Ardavan Salehi Nobandegani
(ardavan.salehinobandegani@mail.mcgill.ca)
Department of Psychology & School of Computer Science, McGill University
Montreal, QC. H3A 1G1 Canada

Paul Thagard (pthagard@uwaterloo.ca)
Department of Philosophy, University of Waterloo
Waterloo, ON. N2L 3G1 Canada

Keywords: artificial intelligence, human intelligence, problem solving, learning, cognitive architecture

Introduction

The past decade has seen remarkable progress in artificial intelligence, with such advances as self-driving cars, IBM Watson, AlphaGo, Google Translate, face recognition, speech recognition, virtual assistants, and recommender systems. Ray Kurzweil and others think that it is only a matter of decades before AI surpasses human intelligence. This symposium will evaluate the extent to which AI currently approximates the full range of human intellectual abilities, and critically discuss the prospects for closing the gap between artificial and human intelligence. Participants will combine the perspectives of computer science, psychology, and philosophy.

The Comparative Cognition of Humans and Machines
Ken Forbus and Dedre Gentner

While there has been great progress in both cognitive science and artificial intelligence, both would benefit from better communication between them. The comparative study of cognition in humans and intelligent machines can shed light on both kinds of systems. In the last decade, the confluence of massive computational resources, massive data sets, and several decades of incremental advances has led to a substantial increase in the ability to build applications with neural networks. Deep learning systems have shown impressive performance in image classification and game learning. However, they still fall far short of capturing human abilities such as explanation and inference, and they require orders of magnitude more data than humans do. We argue that a fundamental lack in these systems is their lack of explicit relational representations. The ability to represent and reason about relational patterns is central to our human ability to explain and predict, and to learn rapidly via analogies with prior knowledge. Fortunately, many of the same factors that have led to gains in deep learning systems are also acting to increase our ability to build large-scale systems with relational representations, which reason and learn in human-like ways. We discuss examples from recent experiments in which analogical learning over relational representations leads to far more humanlike and data-efficient learning than deep learning.

AI and Cognitive Architecture
John E. Laird

There is more talk than ever about general AI, but all the emphasis appears to be on recognition, classification, or reactive decision making with very little on cognition. The emphasis seems to be on only slices of System 1. Within those slices, we see human-level or even super-human performance, but these are very thin slices. Each system is focused on one phenomenon, and given the emphasis on learning from large data sets; it leads to overfitting, not necessarily to specific data, but to the specific problem to the exclusion of developing anything that can work on another problem, or even interact with another cognitive capabilities. In contrast, humans are defined by their flexibility – they can work on many different problems, switching effortlessly from one task to another. They also can learn from many sources of knowledge, on line and in real time, and using a variety of learning techniques. Moreover, they can learn new tasks from scratch in real-time from natural language instruction. A growing field

called Interactive Task Learning has developed an AI system that is embodied in a variety of robotic platforms and that can learn over 50 games and puzzles as well as navigation tasks. It integrates natural language processing, planning, perception, motor control, and learning within a cognitive architecture. Christian Lebiere, Paul Rosenbloom and I have proposed the Common Model of Cognition (CMC) to unify the theoretical underpinnings of many cognitive architectures, starting with Soar, ACT-R, and Sigma. CMC has a vastly different structure than current AI approaches, including procedural and declarative memories, working memory, multiple learning mechanisms. Although these components are common in cognitive science, they are the exception in current AI systems, in large part because of the emphasis on System 1, and off-line batch learning. Until AI takes cognitive architecture, as exemplified by the CMC, seriously, it will not achieve the flexibility, breadth, and adaptability we associate with human intelligence.

Close the Gap and Cooperate

Thomas Shultz and Ardavan Salehi Nobandegani

We will argue that attempts towards achieving artificial general intelligence (AGI) should pay more attention to human intelligence and its neural underpinnings. Having to interact with humans, AGI will need an adequate grasp of human judgment and decision-making and moral principles. Human intelligence not only surpasses current AGI systems, but, importantly, it does so in a resource-efficient way, setting a gold standard for future AI systems. Many of the important AI algorithms originated in psychology, and that strategy is still worth pursuing. A current shortcoming of many AI systems is their limited capacity for generalization – the ability to transfer knowledge from a newly or previously learned task to other relevant tasks. AI could also benefit tremendously from cognitive and developmental psychology to better understand the developmental stages that human infants go through on their way toward adult-level intelligence. To illustrate, we'll focus a bit on the significance of autonomous learning (aka active learning) for bridging the current gap with humans. Even infants take an active role in their own learning by selecting what to work on, what to abandon, and perhaps which examples would be most useful. There is a key role here for learning cessation, the ability to give up on impossible learning tasks, identifiable by lack of continued progress. This paves the way for focusing on tasks in which progress and mastery are more likely. We can suggest ways of implementing these important human capacities in future AI systems. Finally, we want to stress the importance of a cooperative relationship between humans and machines. The notion of gap between us and them that can be closed or even surpassed suggests a more competitive relationship than there perhaps needs to be. The results of mutual cooperation between humans and machines could be much more interesting and desirable to achieve.

How AI Can Understand Causality

Paul Thagard

Causality is important for operating in the world and explaining how it works. Yoshua Bengio and others have pointed out that deep learning and other AI systems lack a human-level understanding of causality. Thagard (2019) argues that human understanding of causality originates with sensory-motor-sensory schemas found in infants as young as 2.5 months. For example, a baby can see a rattle, hit it with hands, and see the rattle move and make a noise. Learning robots could potentially form such schemas, but would have to go beyond current AI systems in several ways. First, they would need *modal retention*, the capacity to save and work with sensory and motor representations. This capacity is found in the Semantic Pointer Architecture of Chris Eliasmith (2013), but not in other cognitive architectures or AI systems. Second, they would need the capacity to learn dynamic patterns that capture changes in series of events. Third, they would need to be able to expand the rudimentary sensory-motor appreciation of causality to cover advanced elements that included regularities, probabilities, and manipulations.

References

Eliasmith, C. (2013). *How to build a brain: A neural architecture for biological cognition*. Oxford: Oxford University Press.

Forbus, K. D., & Gentner, D. (2017). Evidence from machines that learn and think like people. *Behavioral and Brain Sciences*, 40.

Forbus, K., Liang, C. and Rabkina, I. (2017). Representation and computation in cognitive models. *Topics in Cognitive Science*, doi:10.1111/tops.12277.

Laird, J. E., Gluck, K., Anderson, J., Forbus, K., Jenkins, O., Lebiere, C., Salvucci, D., Scheutz, M., Thomaz, A., Trafton, G., Wray, R. E., Mohan, S., Kirk, J. R. (2017). Interactive task learning, *IEEE Intelligent Systems,* 32(4), 6-2.

Laird, J. E., Lebiere, C., & Rosenbloom, P. S. (2017). A standard model of the mind: Toward a common computational framework across artificial intelligence, cognitive science, neuroscience, and robotics. *AI Magazine*, 38(4), 13-26.

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(e253), 1–72.

Shultz, T. R., & Doty, E. (2014). Knowing when to quit on unlearnable problems: another step towards autonomous learning. In J. Mayor & P. Gomez (Ed.), *Computational Models of Cognitive Processes* (pp. 211–221). London: World Scientific.

Thagard, P. (2019). *Brain-mind: From neurons to consciousness and creativity*. New York: Oxford University Press.