

Invited keynote talk:

A neural architecture for reasoning, decision-making, and episodic memory: Taking a cue from the brain

Lokendra Shastri
International Computer Science Institute
Berkeley, CA 94704, USA
www.icsi.berkeley.edu/~shastri

ABSTRACT

The human brain is capable of encoding a large and diverse body of common sense knowledge and performing a wide range of inferences rapidly; to wit our ability to understand language in real-time. Consider the simple narrative: “John fell in the hallway. Tom had cleaned it. He got hurt.” Upon reading this narrative most of us would effortlessly infer that *Tom had cleaned the hallway, therefore, the hallway floor was wet; John slipped and fell because the floor was wet; John got hurt because of the fall*. Such “bridging” inferences help in establishing referential and causal coherence and are essential for language understanding. Since we understand language at the rate of *several hundred words per minute*, it follows that we draw such inferences within hundreds of milliseconds.

The ability to reason effectively with a large body of knowledge is critical not only for language understanding, but also for decision-making, problem solving, and planning. Consequently, the development of efficient, large-scale inference systems has been a central goal of artificial intelligence. Although notable progress has been made toward this goal, it remains far from being achieved.

Given that the human brain is the only extant system capable of supporting a broad range of efficient, large-scale reasoning, it seems appropriate to assume that understanding how the brain represents knowledge and performs inferences might lead to critical insights into the structure and design of large-scale inference systems.

In this talk I will review the current state of a long-term research project on understanding the neural basis of knowledge representation, reasoning, and memory, with an emphasis on the representation and processing of relational (first-order) knowledge. In particular, I will describe SHRUTI, a neurally plausible model that demonstrates how a suitably structured network of simple nodes and links can encode several hundred thousand episodic and semantic facts, causal rules, entities, types, and utilities and yet perform a wide range of explanatory and predictive inferences within a few hundred milliseconds. I will examine some of the predictions stemming from this work about the characteristics of common sense reasoning, and I will discuss how insights arising from this work can be leveraged to design a scalable inference system running on conventional computers.