

Building and Evaluating Models of Human-Level Intelligence

Kenneth Forbus (forbus@northwestern.edu)

Department of Computer Science, Northwestern
2133 Sheridan Road Evanston, IL, 60208

Joshua Tenenbaum (jbt@mit.edu)

Department of Brain and Cognitive Sciences, MIT
77 Massachusetts Ave, Cambridge, MA 02139

Nicholas Cassimatis

Department of Cognitive Science, RPI
110 8th St. Troy, NY 12180

Discussant: **Paul Bello**

Air Force Research Laboratory
32 Brooks Rd., Rome, NY 13441

Motivation

The goal of modeling human-level intelligence has recently received increased attention in cognitive science and artificial intelligence. How do we build and evaluate such models?

What are the main questions in understanding human-level intelligence that modelers must address? It is sufficient in many subfields of cognitive modeling to demonstrate that a model fits and predicts experimental data and is therefore likely to at least approximate the actual mechanisms of cognition. To validate a model of (some aspects of) human level-intelligence, one must also demonstrate that it can behave at a human-level of competence. How do we demonstrate this?

Three aspects of the human-level intelligence problem call into question the traditional role of data in evaluating cognitive models. First, many tasks that have been addressed experimentally (e.g., digit recall, mental rotation) can be executed using several known computational mechanisms. The problem is to decide which of these mechanisms the mind uses. For many tasks requiring human-level intelligence, e.g., natural language dialog and interpreting visual scenes, there are no known algorithms as capable as humans. Is it therefore premature to conduct experiments to choose from among competing models without first producing at least one model that exhibits the abilities humans do?

Second, even once a model is developed that behaves at or near to human-level ability, it is difficult to employ traditional model fit statistics to produce evidence that the mind reflects real mechanisms. Many problems requiring human-level intelligence are dependent on the complex dynamics between many different cognitive processes. Does it therefore make sense to use traditional model fit metrics to evaluate theories of human-level intelligence?

Finally, a hallmark of human-level intelligence is that ability to adapt to and perform well in many situations and tasks. Most experimental designs are specifically designed to measure behavior in one or a few task conditions with minimal variation between them. Do we need new methods to test and quantify the generality and ability to adapt of human intelligence and confirm that our models are powerful enough to explain it?

Kenneth Forbus

Ultimately, models of cognitive processes need to be capable of human-level performance. This requires building simulations at a larger scale than is common today. I think the *integration constraint* is important methodologically: A cognitive simulation of a process should be able to serve as a component in simulations of larger-scale cognitive processes. I will argue that large scale experiments at the level of cognitive architectures and component-level experiments can provide valuable mutual constraint, using examples from our group's work.

Joshua Tenenbaum

Models of human cognition ideally should be judged on at least two criteria: their ability to explain -- or even produce -- human-level competence on an array of challenging tasks, and their ability to make strong testable empirical predictions about behavior, with a minimum of free parameters or ad-hoc auxiliary assumptions. These goals are often treated as incompatible, at least implicitly, with AI researchers typically emphasizing the former and cognitive psychologists emphasizing the latter. I will argue that both goals are necessary for successful cognitive modeling and that it is a practical and productive research strategy to attempt to meet both simultaneously. I will illustrate with some examples from work by my group and others.

Nicholas Cassimatis

Models of cognition are typically evaluated in terms of how well they fit and/or predict certain data sets. I will explore and illustrate from current research several other ways in which models and data can be used in cognitive science. These include: 1. demonstrating that a model has the computational power to explain human-level intelligence; 2. narrowing the inferences we make from subject behavior in a particular task (especially in the developmental literature), 3. demonstrating that cognition in domains can be conducted using the same basic set of mechanisms.

Pre-conference online discussion.

Since these issues are fundamental to many branches of cognitive science, and relatively unexplored, an online discussion forum (e.g., a Wiki) would be set up to discuss these issues in the months preceding the symposium.