MIT Open Access Articles

*The New Frontier of Human-Level
Artificial Intelligence Introduction*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Beal, Jacob, and Patrick H. Winston. "Guest Editors' Introduction: The New Frontier of Human-Level Artificial Intelligence." *Intelligent Systems*, IEEE 24.4 (2009): 21-23. © 2009 IEEE

As Published: <http://dx.doi.org/10.1109/MIS.2009.75>

Publisher: Institute of Electrical and Electronics Engineers

Persistent URL: <http://hdl.handle.net/1721.1/52357>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.

The New Frontier of Human-Level Artificial Intelligence

Jacob Beal,
BBN Technologies

Patrick Winston,
Massachusetts
Institute
of Technology

Within the field of human-level intelligence, researchers are combining a variety of approaches toward the goals of human-like breadth, flexibility, and resilience for artificial intelligence systems.

Humans are still the “gold standard” of intelligent systems. Although machines have surpassed our capabilities in many particular domains, such as solving calculus problems and finding the shortest routes through graphs, no artificial system even comes close to the breadth, flexibility, and integration of capabilities exhibited by the average human. Even in those domains where we generally

regard machines as having attained human-equivalent capability, this remains true only so long as we narrowly limit the domain: while a machine can generally solve calculus problems that a human can't, only a human is capable of sorting out which calculus problems are worth solving, or what a game of chess might reveal about the opponent's personality.

This fact is an inspiration, and not a discouragement. Looking backward, the field has had great success already in attaining human-level capabilities in narrow domains and shedding light on particular aspects of cognition. At the same time, we see the emergence of a new frontier in human-level AI research, where the problems of breadth, flexibility, and integration are beginning to be tackled directly. This yields a prospect both for revolutionary change in the capabilities of machines and also for a synthesis of cognitive models towards a broader understanding of the nature of human cognition.

Why We Aim High

Human-level AI is unapologetic in its lofty goals, which tend to emerge from two dominant and intertwined motivations: better understanding of human intelligence and increasing the capabilities of machines. This field stands squarely at the intersection of cognitive science and computer science. Researchers both draw on the tools of computing and derive important design constraints from knowledge about human and animal intelligence. While these interests sometimes compete, and many human-level AI researchers have at times been left to feel orphaned, more often this duality produces fruitful results for both fields.

On the one hand, the construction of machine systems can aid our understanding of human intelligence by illuminating particular aspects of human cognition. When a cognitive model is realized as a computational process, it's forced to be fully instantiated, often exposing subtle errors, unexpected constraints, and computational challenges. This is a double-edged sword, because it's usually the case that many details of that instantiation can't be grounded in our knowl-

edge of human cognition. Even so, the insights gained from working with an operational model can produce surprising insights and generate new targets for investigation by cognitive psychologists and neuroscientists.

On the other hand, human-like cognitive abilities would be extremely valuable in many applications. Here, the study and modeling of human capabilities contributes both candidate mechanisms and engineering chal-

The ease with which
humans solve problems
that have previously
appeared intractable
is a gauntlet thrown
down to us as engineers.

lenges. The ease with which humans solve problems that have previously appeared intractable, such as learning the meaning of words or reasoning about the beliefs of others, is a gauntlet thrown down to us as engineers, and we often discover new algorithms and engineering principles through the effort to model human capabilities. Though particular models of human-level behavior may work too hard at being faithful to the original to find immediate application, the insights that they yield have caused pragmatic revolutions before (for example, mathematical solvers, CAD tools, relational databases), and will continue to do so.

The New Frontier

The focus of this special issue of *IEEE Intelligent Systems* is on an emerging new frontier in human-level

AI research, where we are beginning to see the convergence of a variety of different approaches toward the goals of human-like breadth, flexibility, and resilience for AI systems.

The research on this new frontier is partially driven by both the continued rapid advances in cognitive studies of human intelligence and by the continued Moore's law increase of computational resources. Just as important, however, are difficulties that have been uncovered in the course of recent large-scale efforts like statistical natural language processing, the construction of the Cyc common-sense knowledge base, and the continuing development of traditional cognitive architectures such as Soar and ACT-R.

We thus see that the base goals of breadth and flexibility seem to be driving human-level AI efforts toward addressing three key challenges:

- scaling to massive amounts of background knowledge and/or computational mechanisms,
- principles that enable multiple representations or independent cognitive systems to cooperate in reasoning, and
- the ability to tolerate incorrectness and contradictions in a system's knowledge or models of the world.

In this special issue, we present four papers at the forefront of the new frontier in human-level AI research. Each brings a different background and perspective on the subject, and hence a different technical approach.

We begin with "Applying Common Sense Using Dimensionality Reduction," where Catherine Havasi, Robert Speer, James Pustejovsky, and Henry Lieberman grapple with scaling and the problems of integrating large data sets. The authors present a dimensionality-reduced represen-

tation of semantic networks called AnalogySpace, which they apply to find patterns, smooth out noise, and predict new knowledge based on the hundreds of thousands of relations in the Open Mind Common Sense ConceptNet. They further extend this to blend ConceptNet together with ontological knowledge from WordNet, generating useful new knowledge despite the incompatibilities between the representations of these two massive data sets and their internal inconsistencies.

Ken Forbus, Matt Klenk, and Tom Hinrichs take a different approach to similar problems in the second article, "Companion Cognitive Systems: Design Goals and Lessons Learned So Far." Here the main concern is integration of knowledge across domains, and the authors are engaged in a research program to see how much can be accounted for by analogical reasoning alone. This necessarily engages them with problems of inconsistencies between models as well as the management of the millions or billions of facts that an intelligent system might incrementally acquire during a lifetime of experience.

In "Reference Resolution Challenges for an Intelligent Agent: The Need for Knowledge," Marjorie McShane presents a careful examination of the range of challenges that an intelligent system faces in determining what objects are referred to by linguistic utterances. Although humans handle all of these with ease, purely syntactic approaches to language lack the necessary information. The article then goes on to lay out a program for addressing these problems through the systematic incorporation of certain categories of knowledge.

Finally, Nicholas Cassimatis examines logical mechanisms for doing just such integration of disparate knowledge sources in "Flexible In-

THE AUTHORS

Jacob Beal is a scientist with BBN Technologies. His research interests center on the engineering of robust adaptive systems, with a focus on problems of system integration for human-level intelligence and on problems of modeling and control for spatially-distributed systems like sensor networks, robotic swarms, and cells during morphogenesis. Beal received his PhD in electrical engineering and computer science from the Massachusetts Institute of Technology. He is a member of the IEEE, AAAI, ACM, and the Cognitive Science Society. Contact him at jakebeal@bbn.com.

Patrick H. Winston is the Ford Professor of Artificial Intelligence and Computer Science at the Massachusetts Institute of Technology (MIT). His Genesis Group studies how vision, language, and motor faculties account for intelligence, integrating work from several allied fields, including not only AI, but also computer science, systems neuroscience, cognitive science, and linguistics. Winston received his PhD in Computer Science from MIT. He is a past president of AAAI and a member of the Naval Research Advisory Committee. Contact him at phw@cmail.mit.edu.

ference with Structured Knowledge through Reasoned Unification." Operating within the model of a cognitive substrate, the article develops a mechanism called "reasoned unification" that fills in missing information in script and frame representations by posing questions about identity relations to one or more reasoning mechanisms. This mechanism provides a rational approach to how different cognitive capabilities can work together to interpret ambiguous, implied, and nonliteral references.

These four articles are united by

a common thread: each addresses a problem in achieving human-like breadth of capability, and is thus led to engage with problems of scaling and integration. Most importantly, however, each tells a clear story of how progress on the specific problems the authors are working on today leads to progress on the larger investigation of human-level intelligence. ■

For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.