

Report on the First Conference on Artificial General Intelligence (AGI-08)

Hugo de Garis and Ben Goertzel

- *The First Conference on Artificial General Intelligence (AGI-08) was held on March 1–3, 2008, at the University of Memphis. The overall goal of the conference was to work toward a common understanding of the most promising paths toward creating AI systems with general intelligence at the human level and beyond and to share interim results and ideas achieved by researchers actively working toward powerful artificial general intelligence.*

When the AI field began in the middle of the last century, its main focus was on the creation of computer programs displaying intelligence at the human level and ultimately beyond. This proved harder than anticipated, and over the following decades most research activity in the field shifted toward the creation of computer programs solving specific sorts of problems that humans find to require intelligence. This more task-focused “narrow AI” work that has come to dominate the field has yielded a variety of theoretical and practical successes, but there is no real consensus regarding the relation of this work with the original goal of creating powerful general intelligence at the human level and beyond. In the last few years, an increasing number of AI researchers have come to feel the time is ripe for a renewed focus on the original goals of the AI field; and this was the motivation for the convening of the First Conference on Artificial General Intelligence (AGI-08), which was held at the FedEx Institute at University of Memphis in early March 2008. Detailed information about the conference, including a complete proceedings, is available online for those who were not able to attend.¹

The conference was characterized by an enthusiastic and energetic tone, and there was a feel of real excitement about having so many AGI-oriented researchers in one place at one time. There were more than 120 attendees (including roughly 40 presenters). The participants included not only AGI researchers from academe and industry,

but also scientists and engineers from allied disciplines and a smattering of artists, business people, and other AGI enthusiasts, who provided a valued breadth to the discussions. The crowd was international, with nearly as many European as American participants.

The conference room was a futuristic setting called “The Zone,” with a vague resemblance to the Star Trek bridge, including an excellent if mildly glitchy video system that, during Q&A sessions, displayed the questioner on a big screen in front of the room. This worked well with the unconventional format of the conference, which was divided into topical sessions, each one featuring 10-minute summary talks by researchers, followed by a lengthy moderated discussion session. Some of the sessions were markedly interdisciplinary; for instance, the session on AI in Virtual Worlds was chaired by Sibley Verbeck (CEO of Electric Sheep Company); and the session on neural nets was chaired by Randal Koene (a neuroscientist from Boston University).

A postconference workshop on “Sociocultural, Ethical, and Futurological Implications of AGI,”² which drew about 60 participants, was introduced by noted futurist Natasha Vita-More and included some wide-ranging and passionate discussions.

Why Might the Time Be Ripe for AGI?

Given the checkered history of ambitious attempts at human-level AGI, it may justly be wondered whether the time is yet ripe for another volley of such attempts. This was one of the main topics of discussion in the opening session of the conference. Following a general introduction by conference chair Stan Franklin (University of Memphis), Ben Goertzel gave a talk on “The Past, Present, and Future of Artificial General Intelligence.” While the researchers in the audience displayed a healthy level of disagreement regarding this topic, the overall opinion was that it does make sense, at present, to focus considerable resources and energy directly on the AGI problem, as distinct from more task-spe-

cific, narrowly-focused AI initiatives.

As an example of the issues raised in the discussion following Goertzel’s introductory talk, Hugo de Garis asked Goertzel: “People have been trying for 60 years to produce human-level AI, and have failed. Why do you think the time is ripe to try again? What has changed, if anything?”

Goertzel responded: “I think four things have changed: (1) Computers are far faster now, and far more easily networkable. (2) Understanding of cognitive science and cognitive neuroscience has advanced sufficiently that we can make an empirically grounded “high-level architecture diagram” for the human mind, sufficient to guide the development of humanlike minds. (3) Virtual world and robotics technology have advanced dramatically so that it’s far more feasible to experiment with virtually and physically embodied AI. (4) Narrow-AI has advanced algorithmically hugely, for instance we can now solve Boolean satisfaction problems with hundreds of thousands of variables. We can use automated program learning methods to automatically learn $n \cdot \log(n)$ sorting algorithms based on experience

Following a review of these and other general issues in the introductory session, the main body of the conference focused on specific ideas regarding the construction of AGI systems and preliminary results obtained by creators of present-day proto-AGI systems, including well-known systems such as SOAR and OpenCyc, as well as systems of more recent creation based on more novel approaches.

Presentations and Discussions

The session topic headings of the conference reflected closely those areas of AGI that the conference program committee (led by Ben Goertzel, Pei Wang, and Stan Franklin) felt critical to future progress in AGI. There were seven such sessions. There was also a discussion session on the coming AGI renaissance and a business meeting, at which AGI-09 was proposed for the following year, to be held very probably in Washington, DC. It was also de-

cidated to launch *JAGI (Journal of Artificial General Intelligence)*, a new open source journal with Pei Wang as chief executive editor and Hugo de Garis, Ben Goertzel, and Kai-Uwe Kühnberger as coexecutive editors.³

The highlights of the scientific presentations were numerous; however, the nonresearchers and media representatives in attendance seemed most energized by the session on AI in Virtual Worlds, perhaps because in this session the presenters (Goertzel, Andrew Shilliday, and Martin Magnusson) showed movies of animated characters carrying out intelligent behaviors.

Shilliday’s presentation of work done by Selmer Bringsjord’s team at RPI resulted in a considerable degree of media attention, with headlines such as “AI Program Thinks Like Four Year Old”⁴ and “Childlike Intelligence Created in Second Life.”⁵ On a technical level, the work involved using a logic-based AI system to control a humanoid virtual agent in the Second Life virtual world, which interacted with other agents in a manner displaying what developmental psychologists call “Theory of Mind,” an ability that generally emerges in humans at around four years of age.

Ben Goertzel’s presentation in this session also discussed AI in Second Life, but from a different perspective: the creation of virtual dogs that learn through imitation, reinforcement, and correction.

The session on AGI Architectures was also a particularly vibrant one, including a great presentation by John Laird on the latest incarnation of the SOAR cognitive architecture, which includes a variety of features not present in the classical SOAR design that (in the view of the authors and the majority of AGI-08 discussion participants) palpably move SOAR closer to being a viable architecture for AGI.

One of the main points to come out of the discussion in the AGI Architecture session was the degree of commonality among diverse AGI designs. It seemed clear, in the context of the session, that several separate research programs are pursuing fairly similar architectures and ideas but using different languages to describe what

they're doing. This suggests that making a systematic effort at finding a common language and really understanding the true overlaps and differences of the various approaches might be very beneficial. There was some talk of organizing a small, invitation-only workshop among practicing AGI system architects, perhaps in fall 2008, with a view toward making progress in this direction.⁶

Along similar lines, in a later session of the conference was also discussed an "AGI Roadmap" project that would involve aligning different cognitive architectures currently proposed insofar as possible but also go beyond that. For instance, another key aspect of the roadmap might be an agreement on certain test environments or tasks that could be used to compare and explore various AGI architectures in more of a common way than is now possible.

Final Remarks

Society, including the society of scientists, is starting to wake up to the notion that, given recent advances in various areas of technology and science, human-level AGI is not necessarily a pipe dream but might be rationally considered a potential near-term reality. There is still a hard path to tread before the core task of figuring out how to create powerful AGI software is solved; but the collective energy felt at AGI-08 was highly inspiring for those of us who believe the time is ripe for the community to make a renewed and coordinated attack on the problem.

Notes

1. See agi-08.org
2. See agi-08.org/workshop/
3. See journal.agi-network.org
4. See www.itnews.com.au/News/72057_childlike-intelligence-created-in-second-life.aspx
5. See www.technewsworld.com/rsstory/62073.html
6. In part, as a result of this discussion at AGI-08, a workshop was held in October 2008 at the University of Michigan. It focused on evaluation and metrics for human-level AI.

Hugo de Garis is director of the "China

CAMBRIDGE

New and Forthcoming Titles from Cambridge University Press!

Forthcoming...
Statistical Machine Translation
 Philipp Koehn
 \$60.00: Hardback: 978-0-521-87415-1: 488 pp.

New!
Handbook of Practical Logic and Automated Reasoning
 John Harrison
 \$150.00: Hardback: 978-0-521-89957-4: 704 pp.

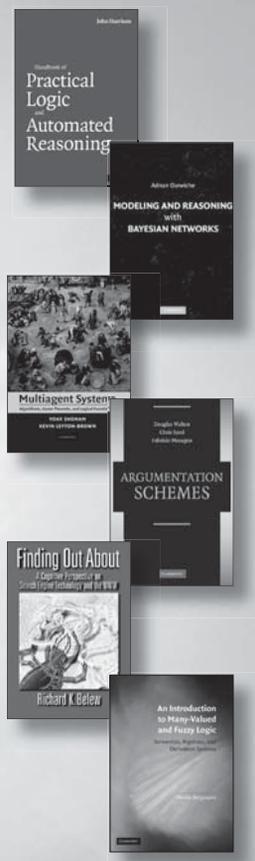
Modeling and Reasoning with Bayesian Networks
 Adnan Darwiche
 \$95.00: Hardback: 978-0-521-88438-9: 576 pp.

Multiagent Systems
 Algorithmic, Game-Theoretic, and Logical Foundations
 Yoav Shoham and Kevin Leyton-Brown
 \$60.00: Hardback: 978-0-521-89943-7: 496 pp.

Argumentation Schemes
 Douglas Walton, Christopher Reed, and Fabrizio Macagno
 \$90.00: Hardback: 978-0-521-89790-7: 436 pp.
 \$32.99: Paperback: 978-0-521-72374-9

New in Paperback!
Finding Out About
 A Cognitive Perspective on Search Engine Technology and the WWW
 Richard K. Belew
 \$45.00: Paperback: 978-0-521-73446-2: 384 pp.

An Introduction to Many-Valued and Fuzzy Logic
 Semantics, Algebras, and Derivation Systems
 Merrie Bergmann
 \$90.00: Hardback: 978-0-521-88128-9: 342 pp.
 \$36.99: Paperback: 978-0-521-70757-2



Prices subject to change.

www.cambridge.org/us/computerscience
 1-800-872-7423



CAMBRIDGE
UNIVERSITY PRESS

Brain Project," which is attempting to build China's first artificial brain by evolving tens of thousands of neural net modules quickly using Celoxica accelerator boards, downloading the results into the memory of a PC, connecting them to make an artificial brain, and having the PC perform the neural signaling of the whole brain in real time to control the hundreds of behaviors of robots. He is the author of "The Artelect War," which discusses the politics of the rise of massively intelligent machines later this century. He is currently professor of computer science and mathematical physics at Xiamen University, China.

Ben Goertzel is CEO of AI companies Novamente LLC and Biomind LLC and the leader of the Novamente Cognition Engine and OpenCog AGI projects. He is the author of seven and editor of four books on artificial intelligence and related topics, along with more than 70 research papers and numerous journalistic articles. He has also led teams developing a number of AI-based commercial software products in bioinformatics, natural language processing, virtual agent control, and other areas. He is also director of research at the non-profit Singularity Institute for AI.