

# The Second International Conference on Artificial General Intelligence (AGI-09)

*Hugo de Garis and Ben Goertzel*

■ This article reports on AGI-09, the Second International Conference on Artificial General Intelligence, which was held 6–9 March 2009 in Arlington, Virginia, USA. The purpose of the conference was to exchange ideas about the creation of artificial systems with general intelligence at, and ultimately beyond, the human level.

AGI-09, the Second International Conference on Artificial General Intelligence, was held March 6–9 in Arlington, Virginia. Ben Goertzel chaired the conference, and Marcus Hutter and Pascal Hitzler chaired the program committee. Continuing the mission of AGI-08 (which was held March 2008 at the University of Memphis), the purpose of the conference was to provide a venue for exchange of in-depth scientific ideas and results between researchers working directly toward the original goal of the AI endeavor: the creation of artificial systems with general intelligence at the human level and ultimately beyond.

The first day of the conference featured in-depth tutorials on leading AGI systems and approaches, including introductions to the SOAR, Texai, and OpenCog software, and overviews of the logic-based, reinforcement learning and program-induction approaches to AGI. Following this, the main conference on Saturday and Sunday featured a number of themed sessions: Evaluation and Metrics (chaired by John Laird), Robotics and Embodiment (chaired by Itamar Arel), Cognitive Architectures (chaired by Pei Wang and Stephen Reed), Logical Approaches to AGI (chaired by Selmer Bringsjord), Learning and Reasoning (chaired by Selmer Bringsjord), Speech and Language (chaired by Moshe Looks), and Self-Awareness and Consciousness (chaired by Ben Goertzel).

As compared to its predecessor, AGI-09 generally had more of a sophisticated, theoretical-computer-science-based tone. There were fewer industry participants because in early 2009 (due to the global economic crisis) many U.S. firms were restricting

employee travel. On the other hand there was an even greater international participation, including a keynote speech by Juergen Schmidhuber (from IDSIA, in Lugano, Switzerland, and the Technical University of Munich) and a large number of presentations from German researchers.

In his keynote, Schmidhuber noted that since age 15 or so his main scientific ambition has been “to build an optimal scientist, then retire.” After an impressive overview of the work his team has done with recurrent neural net architectures, and a discussion of ideas from the mathematical theory of general intelligence, Schmidhuber wrapped up his talk with a futurist thematic, presenting a lively argument in favor of the hypothesis that technological development is propelling the human race toward a singularity: a point after which development on Earth is driven by AIs with vastly greater intelligence than humans, so that humans no longer have control or even comprehension of the progress of events. This hypothesis is best known in the versions proposed by Vernor Vinge (1993) and Ray Kurzweil (2005)—and the latter has extrapolated recent technology trends to come up with a projected singularity date of 2045. Schmidhuber’s different but related extrapolations lead to a date of 2040. Ever the ironist, however, Schmidhuber followed up this prognostication with another extrapolation, analyzing historical data so as to prove that the Singularity should have occurred in 1540 AD and concluding that, whether the singularity arrives or not, “we are fortunate to live in very interesting times.”

As appropriate for a field in the exploratory stage of development, the more than 40 technical papers presented during the conference were highly diverse, without any single unifying theme. However, there seemed to be a particularly large number of high-quality papers dealing with the problem of “program learning”—that is, the creation of software that can take as input a certain behavior (presented as a description, or as a set of examples), and produce as output a computer program capable of demonstrating this behavior. As Bill Hibbard from the University of Wisconsin noted, “The idea of program learning is not new, but it was remarkable at AGI-09 the extent to which this approach has become the consensus of the best people working in AGI. The presentations of Moshe Looks, Eric Baum, Marcus Hutter, Josh Hall, Juergen Schmidhuber, and several others all focused on systems or methods for learning new algorithms.”

Ray Kurzweil, one of several AGI-09 sponsors, sponsored the “Kurzweil Best AGI Paper Prize.” It provided a \$1000 check to the winner and \$100 checks to two runners-up. The winning paper was titled “Combining Analytical and Evolutionary Inductive Programming,” by Neil Crossley,

Emanuel Kitzelmann, Martin Hofmann, and Ute Schmid, from the Cognitive Systems Group at the University of Bamberg, who work in the AI tradition of “inductive programming.” Their paper described a clever way to reformulate the conclusions of analytic program learning, in a way that doesn’t get scrambled by the evolutionary process, but rather gives evolutionary program learning a head start, allowing it to proceed more effectively than if it were operating without coupling to an analytic process.

The final day was devoted to a workshop on the future of AGI, and the presentations and discussions ranged far and wide. There was an economics-focused session, with presentations such as “If AGI does all the jobs, how will people get income?” (James Albus), “The Economics of AI” (Robin Hanson), and “Artilect War or Utopia” (Hugo de Garis). The ethics session covered topics such as “Unethical but Rule-Bound Robots Would Kill Us All” (Selmer Bringsjord) and “Is Sousveillance the Best Path to Ethical AGI and a Safe Singularity?” (Ben Goertzel). Finally, Itamar Arel led a discussion on the need for an “AGI Roadmap” to help coordinate efforts by multiple researchers and aid in evaluation of results and communication of ideas. There was wide agreement that a roadmap initiative could be valuable for moving the AGI field forward, and this is likely to be a focus of effort by a subset of the AGI community over the next couple years.

## References

Kurzweil, Ray 2005. *The Singularity Is Near: When Humans Transcend Biology*. New York: Viking

Vinge, Vernor 1993. The Coming Technological Singularity. In *Vision-21: Interdisciplinary Science and Engineering in the Era of CyberSpace*. NASA Conference Publication CP-10129. Washington, D.C.: National Aeronautics and Space Administration.

**Hugo de Garis** is director of the China 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 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 nonprofit Singularity Institute for AI.