Mapping the Landscape of AGI

Ideas and Conclusions from the 2009 AGI Roadmap Workshop

Ben Goertzel
CEO, Novamente LLC and Biomind LLC
CTO, Genescient Corp
Co-founder, OpenCog Project
Vice Chairman, Humanity+
Adjunct Research Professor, Xiamen University, China
Advisor, Singularity University and Singularity Institute

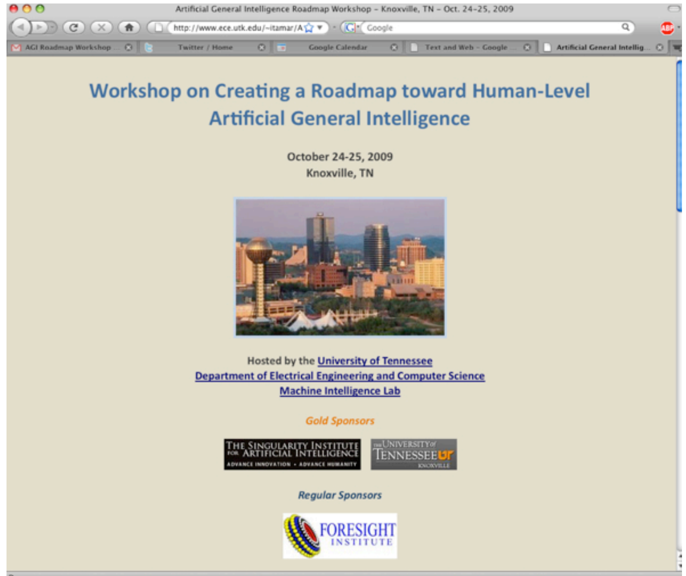
Artificial General Intelligence Roadmap Workshop – Knoxville, TN – Oct. 24-25, 2009
Workshop on Creating a Roadmap toward Human-Level
Artificial General Intelligence
October 24-25, 2009
Knoxville, TN
Hosted by the University of Tennessee
Department of Electrical Engineering and Computer Science
Machine Intelligence Lab
Gold Sponsors
THE SINGULARITY INSTITUTE FOR ARTIFICIAL INTELLIGENCE
ADVANCE INNOVATION • ADVANCE HUMANITY
THE UNIVERSITY of TENNESSEE UT KNOXVILLE
Regular Sponsors
FORESIGHT INSTITUTE

Mapping the Landscape of Human-Level Artificial General Intelligence

Sam S. Adams
IBM Research

Itamar Arel
University of Tennessee, Knoxville

Joscha Bach
Humboldt University

Robert Coop
University of Tennessee, Knoxville

Rod Furlan
Singularity University

Ben Goertzel
Novamente LLC

J. Storrs Hall
Institute for Molecular Manufacturing

Alexei Samsonovich
George Mason University

Matthias Scheutz
University of Indiana

Matt Schlesinger
Southern Illinois University, Carbondale

Stuart C. Shapiro
SUNY Buffalo

John Sowa
Vivomind

Narrow AI


Web Images Videos Maps News Shopping Gmail more ▾
bengoertzel@gmail.com | iGoogle | Settings ▾ | Sign out
Google
Will Google ever be truly intelligent?
Advanced Search
Language Tools
Google Search
I'm Feeling Lucky


03
MDV
Red Bull


Game viewer - Kasparov vs. Deep Blue
File Help
GAME 6
IBM
3.5
2.5
Deep Blue
Kasparov
white
black
MAURICE ASHLEY: After Ngf6 Deep Blue has responded instantly with Bf1-d3, developing the bishop, putting it on a very solid square. Potentially Kasparov might castle king-side, so the

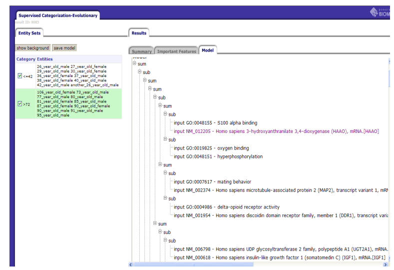

AGI ?

The ability of a system to achieve a variety of complex goals in a variety of complex environments using limited computational resources

(including goals and environments that were not anticipated at the time the system was created)

"I claim achieving real Human-Level artificial intelligence would necessarily imply that most of the tasks that humans perform for pay could be automated. Rather than work toward this goal of automation by building special-purpose systems, I argue for the development of general-purpose, educable systems that can learn and be taught to perform any of the thousands of jobs that humans can perform. Joining others who have made similar proposals, I advocate beginning with a system that has minimal, although extensive, built-in capabilities. These would have to include the ability to improve through learning along with many other abilities."

-- Nils Nilsson,
in the 2005 AI Magazine article
"Human-Level Artificial Intelligence? Be Serious!"

John Laird and Robert Wray (AGI-10): Cognitive Architecture, Environment & Task Requirements for Achieving Human-Level AGI

R0. When given new tasks, structure is not changed via reprogramming [Note: the original R0 was "Fixed structure for all tasks"]

R1. Realize a symbol system

Represent and effectively use:

- R2. Modality-specific knowledge
- R3. Large bodies of diverse knowledge
- R4. Knowledge with different levels of generality
- R5. Diverse levels of knowledge
- R6. Beliefs independent of current perception
- R7. Rich, hierarchical control knowledge
- R8. Meta-cognitive knowledge

R9. Support a spectrum of bounded and unbounded deliberation

R10. Support diverse, comprehensive learning

R11. Support incremental, online learning

C1. Environment is complex with diverse, interacting objects

C2. Environment is dynamic

C3. Task-relevant regularities exist at multiple time scales

C4. Other agents impact performance

C5. Tasks can be complex, diverse and novel

C6. Agent/Environment/Task interactions are complex and limited

C7. Agent computational resources are limited

C8. Agent existence is long-term and continual

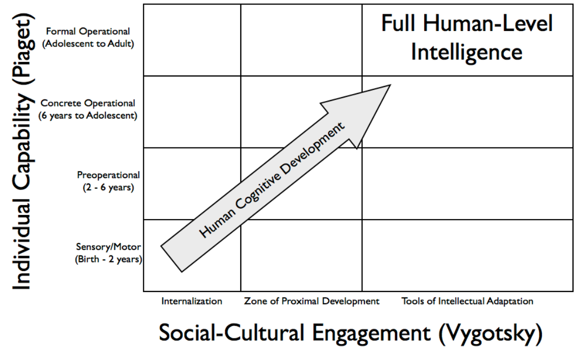
Full Human-Level Intelligence
Human Cognitive Development
Individual Capability (Piaget)
Formal Operational (Adolescent to Adult)
Concrete Operational (6 years to Adolescent)
Preoperational (2 - 6 years)
Sensory/Motor (Birth - 2 years)
Internalization
Zone of Proximal Development
Tools of Intellectual Adaptation
Social-Cultural Engagement (Vygotsky)

Scenarios

- General Video Game Learning
- Preschool Learning (Robotic and/or Virtual)
- Story/Scene Comprehension
- Reading Comprehension
- School Learning
- Wozniak "Coffee Test"

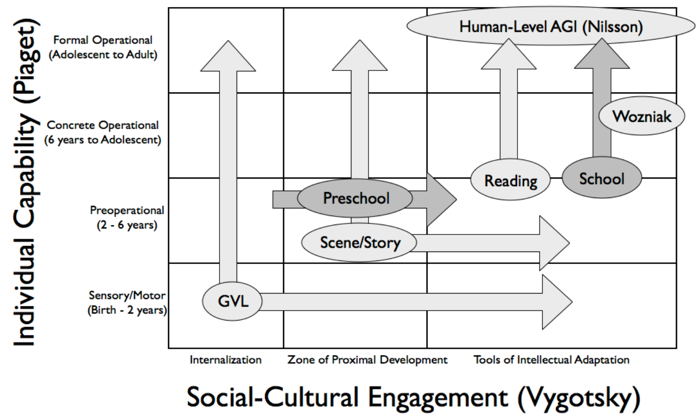

Individual Capability (Piaget)
Formal Operational (Adolescent to Adult)
Concrete Operational (6 years to Adolescent)
Preoperational (2 - 6 years)
Sensory/Motor (Birth - 2 years)
Human-Level AGI (Nilsson)
Wozniak
Reading
School
Preschool
Scene/Story
GVL
Internalization
Zone of Proximal Development
Tools of Intellectual Adaptation
Social-Cultural Engagement (Vygotsky)

Broad competency areas	Sub-areas...	Sub-areas..	Sub-areas..	Sub-areas..	Sub-areas..	Sub-areas..
Perception	Vision	Audition	Touch	Proprioception	Cross-modal	
Actuation	Physical skills	Tool use	Navigation	Proprioception		
Memory	Implicit	Working	Episodic	Semantic	Procedural	
Learning	Imitation	Reinforcement	Dialogical	Via Written Media	Via Experimentation	
Reasoning	Deduction	Induction	Abduction	Causal	Physical	Associational
Planning	Tactical	Strategic	Physical	Social		
Attention	Visual	Social	Behavioral			
Motivation	Subgoal creation	Affect-based				
Emotion	Emotional expression	Understanding emotions	Perceiving emotions	Control of emotions		
Modeling self and other	Self-awareness	Theory of mind	Self-control	Other-awareness	Empathy	
Social interaction	Appropriate behavior	Social communication	Social inference	Cooperation, e.g. group play		
Communication	Gestural	Verbal	Pictorial	Language acquisition	Cross-modal	
Quantitative	Counting observed entities	Grounded small number arithmetic	Comparison of quantitative properties of observed entities	Measurement using simple tools		
Building/ creation	Physical construction w/ objects	Formation of novel concepts	Verbal invention	Social organization		

Scenario	Competency Area	Sub-area	Example Task or Task Family
Virtual Preschool	Learning	Dialogical	**Learn to build a particular structure of blocks** (say, a pyramid) faster based on a combination of imitation, reinforcement and verbal instruction, than by imitation and reinforcement without verbal instruction
Virtual Preschool	Modeling Self and Other	Theory of Mind	While Sam is in the room, Ben puts the red ball in the red box. Then Sam leaves and Ben moves the red ball to the blue box. Sam returns and Ben asks him where the red ball is. **The agent is asked where Sam thinks the ball is.**
Virtual School Student	Learning	via Written Media	Starting from initially available basic concepts (a number, a variable, a function), demonstrate academic progress in **learning how to solve problems from the textbook** using techniques described in the same textbook. The agent should move step by step, from simple to advanced problems, from one domain to another.
Virtual School Student	Modeling Self and Other	Other-Awareness	**Help a friend to cheat on exam** (modified from [REF Samsonovich, 2006]. The virtual agent and a student take an exam consisting of one and the same problem. First, the agent enters a room and works on the problem for an hour, then the student does the same. When the agent exits the exam room after solving the problem, the student comes in, and the agent has a chance to pass a small piece of paper to the student. The entire solution of the exam problem cannot fit on this piece of paper, not even a substantial part of it. The challenge for the agent is to write a short message that would give a useful hint to the friend.
Robot Preschool	Actuation	Proprioception	The teacher moves the robot's body into a certain configuration. The robot is asked to restore its body to an ordinary standing position, and then **repeat the configuration that the teacher moved it into.**

Scenario	Competency Area	Sub-area	Example Task or Task Family
Wozniak Coffee Test	Communication	Gestural	In many cases the robot will be shown to the kitchen. It must understand gestures indicating that it should follow an indicated path, or know how to follow its guide, and know when either is appropriate.
Wozniak Coffee Test	Actuation	Navigation	The robot must go about its business without running into people, walls, furniture, or pets.
Wozniak Coffee Test	Social Interaction	Appropriate behavior	The robot had better be able to recognize the case where it has knocked on the wrong door and the householder is not inviting it in.
Wozniak Coffee Test	Reasoning	Physical	Consider the state of knowledge that tells us we can use a drip pot without its top, but not a percolator. This may come from physical simulation, based on an understanding of naive physics.
Wozniak Coffee Test	Reasoning	Induction	On the other hand, the above-mentioned knowledge about drip pots and percolators may be gathered via inductive reasoning based on observations in multiple relevant situations.

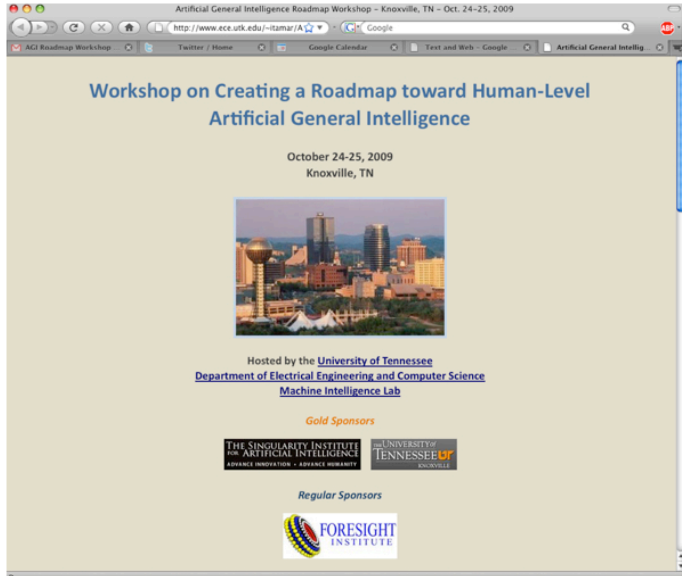
Artificial General Intelligence Roadmap Workshop – Knoxville, TN – Oct. 24-25, 2009
Workshop on Creating a Roadmap toward Human-Level Artificial General Intelligence
October 24-25, 2009
Knoxville, TN
Hosted by the University of Tennessee
Department of Electrical Engineering and Computer Science
Machine Intelligence Lab
Gold Sponsors
THE SINGULARITY INSTITUTE FOR ARTIFICIAL INTELLIGENCE
THE UNIVERSITY of TENNESSEE UT KNOXVILLE
Regular Sponsors
FORESIGHT INSTITUTE