



Project PrimeAGI

Creating a human level or greater **Artificial General Intelligence**

Timing Goal: by **2025 to 2030**

Doing it as an **Open Source** project for the good of all humanity

Doing it **Safely**

Raising funds for Phase 1: 3 years, **\$US 6 million**

Artificial General Intelligence (AGI)

A system capable of choosing actions to achieve a set of complex goals in a variety of complex environments using a realistic amount of computation

The ability to transfer learning from one domain to other domains.

Human Level AGI

- The Turing Test: A text dialogue AI that is indistinguishable from a human in online conversation.
- Better Examples of Human Level AGI:
 - Wozniak's Coffee Test: "when a robot can go into an average American house and figure out how to make coffee, including identifying the coffee machine, figuring out what the buttons do, finding the coffee in the cabinet, etc."
 - Robot University Student Test: "when a robot can enroll in a human university and take classes in the same way as humans, and get its degree, then we've created artificial general intelligence"

What is an **Artificial General Intelligence (AGI)** ?

"a system capable of choosing actions to achieve a set of complex goals in a variety of complex environments using a realistic amount of computation"

"general intelligence is the ability to transfer learning from one domain to other domains."

Human Level AGI

The Turing Test:

A text dialog AI that is indistinguishable from a human in online conversation.

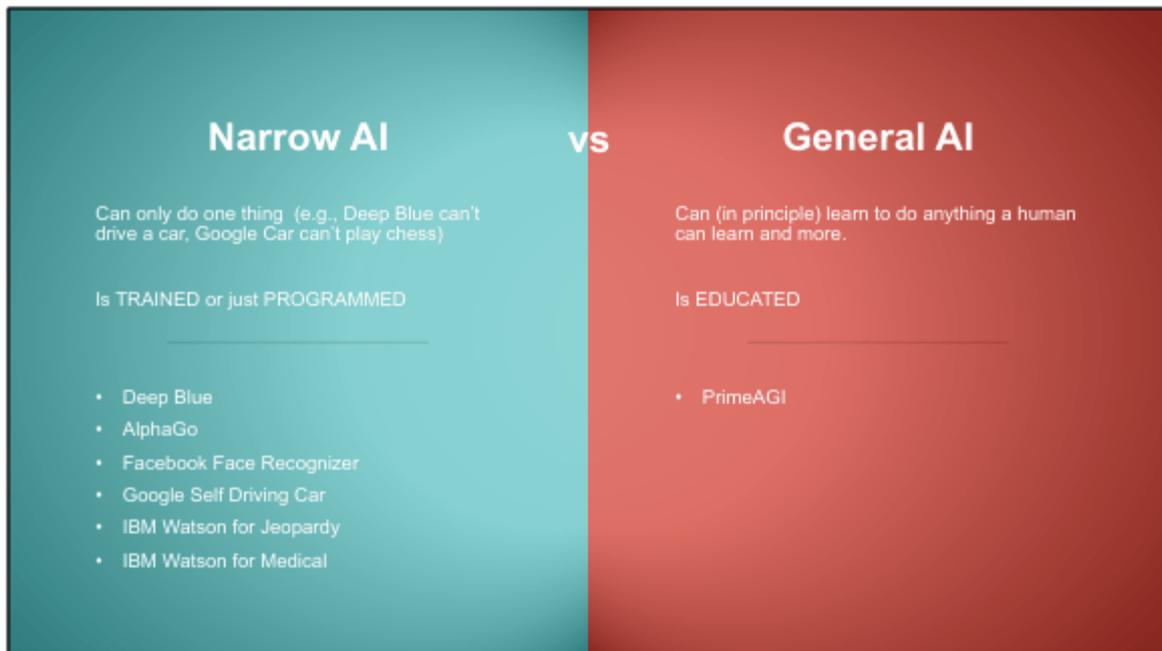
Better Examples of Human Level AGI:

Wozniak's Coffee Test: "when a robot can go into an average American house and figure out how to make coffee, including identifying the coffee machine, figuring out what the buttons do, finding the coffee in the cabinet, etc."

Robot University Student Test: "when a robot can enroll in a human university and take classes in the same way as humans, and get its degree, then we've created

artificial general intelligence”

TOO many references to Robots??



Compare **General** with **Narrow** Artificial Intelligence

Narrow:

Deep Blue
AlphaGo
Facebook Face Recognizer
Google Self Driving Car
IBM Watson for Jeopardy
IBM Watson for Medical

...

A Narrow AI can only do one thing. Eg: Deep Blue can't drive a car, nor Google Car play chess.

A General AI can in principle learn to do anything a human can learn and more.

A Narrow AI is TRAINED or just PROGRAMMED, a General AI is EDUCATED

Why Now?

- AGI has been the implicit long term goal of AI Research since at least 1950
 - Early AI researchers sought human-level AI, but underestimated the problem
 - From the mid-1970s through the mid-aughts, AI research focused on narrow problems,
 - Over the last 10 years, narrow AI heavy application in industry, and serious interest in AGI R&D has resurged
- Today, for the first time, there exists the necessary convergence of research progress and ever increasing computational power and tools to make a serious assault on the AGI problem:
 - Computing hardware and networking have advanced massively
 - Software tools and libraries are far more refined
 - Cognitive science has given us an overview of human mind architecture
 - Narrow AI and CS now provide a formidable arsenal of algorithms and representations
- Others think so too:
 - Kurzweil, Hinton, etc to Google, Andrew Ng to Baidu, Yann LeCun to Facebook, DeepMind acquired by Google for \$500M, Vicarious: \$40M funding from Zuckerberg, GoodAI: \$10M from Marek Rosa directly for human-level AGI, OpenAI – \$1B committed for future work on open-source AI...

Why Now?

AGI has been the implicit long term goal of AI Research since at least 1950

Early AI researchers sought human-level AI, but underestimated the problem
From the mid-1970s through the mid-aughts, AI research focused on narrow problems,

Over the last 10 years, narrow AI heavy application in industry, and serious interest in AGI R&D has resurged

Today, for the first time, there exists the necessary convergence of research progress and ever increasing computational power and tools to make a serious assault on the AGI problem:

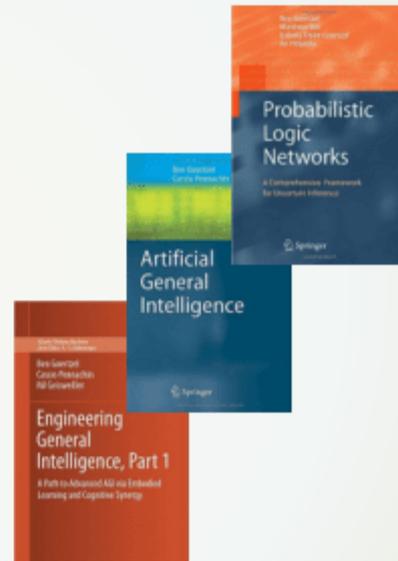
Computing hardware and networking have advanced massively
Software tools and libraries are far more refined
Cognitive science has given us an overview of human mind architecture
Narrow AI and CS now provide a formidable arsenal of algorithms and representations

Others think so too:

Kurzweil, Hinton, etc to Google, Andrew Ng to Baidu, Yann LeCun to Facebook,
DeepMind acquired by Google for \$500M, Vicarious: \$40M funding from Zuckerberg,
GoodAI: \$10M from Marek Rosa directly for human-level AGI, OpenAI – \$1B
committed for future work on open-source AI...

Why Us?

- Ben Goertzel
 - leading thinker, evangelist and organizer in AI community
 - 15 books, 100+ papers
 - Founder and leader of
 - annual AGI Conference
 - Open Source AGI Foundation
 - OpenCog project
 - 2 decades practical AI application experience in robotics, finance, biology, NLP, data mining, etc.
- Core team with 20 year learning curve working together:
 - Webmind -> Novamente -> Opencog -> PrimeAGI
- Opencog is a solid foundation on which to build: significant open source community with 50 current participants plus "alumni association"
- A specific point of view: an approach and a plan – not just a general AGI research effort – Reasonably well documented in *Engineering General Intelligence*, Vols. 1 and 2
- Firm commitment to Open Source



Why Us?

Ben Goertzel

leading thinker, evangelist and organizer in AI community

15 books, 100+ papers

Founder and leader of

annual AGI Conference

Open Source AGI Foundation

OpenCog project

2 decades practical AI application experience in robotics, finance, biology, NLP, data mining, etc.

Core team with 20 year learning curve working together:

Webmind -> Novamente -> Opencog -> PrimeAGI

Opencog is a solid foundation on which to build: significant open source community with 50 current participants plus "alumni association"

A specific point of view: an approach and a plan – not just a general AGI research effort – Reasonably well documented in *Engineering General Intelligence, Vols. 1 and 2*

Firm commitment to **Open Source**

Open Source

- AGI is too important for humanity to be the property of one commercial or governmental entity.
- Creating safe, beneficial AGI is a difficult problem and the best solution is more likely to emerge if
 - a wide variety of bright minds are working together on it
 - It is created within an effort with beneficial AGI as the central goal, rather than e.g. profitability of a particular corporation or military success of a particular country
- High leverage on investment through growing network of volunteers, academics, and commercial Open Source community participants:
 - Directly funding 10 developers may indirectly bring 1000 additional developers funded by other sources



VS



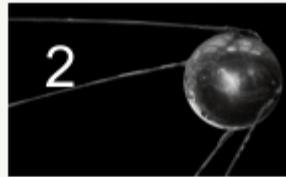
Project PrimeAGI

Use and extend the existing Opencog AGI Framework to drive towards at least Human level AGI in three phases:



"V2 Rocket of AGI"
First system to show exciting functionality that demonstrates the core principles needed to make AGI work

3-5 years



"Sputnik of AGI"
First system to show incontrovertible evidence of AGI

+3-5 years



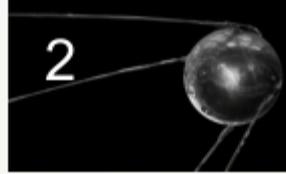
"Apollo of AGI"
Human-level AGI

+3-5 years

9-15 years total



Robot head that holds a simple conversation with a human, displaying genuine understanding of its embodiment and context



Robot head that learns via interaction with humans and the world, similarly to a human child (including learning new languages and social skills via experiential interaction)



Robot head that holds a simple conversation with a human, displaying genuine understanding of its embodiment and context

Once an AGI can be taught university-level mathematics and CS, we begin moving toward profoundly self-improving, self-modifying AGI.

1

“V2 Rocket”

Goal of this fundraising round

A realistic-looking, emotionally and socially interactive robot head which holds a simple conversation with you. The conversation isn't just a “chat bot” but an embodied intelligent dialogue system

Robot dialogue displays genuine understanding of:

- who and what it is, including its own history
- who you are
- what its physical surroundings are

System able to learn and reason about its environment and interactions

- Learning/reasoning ability broadly similar to a young human child
- In some cases more advanced – “idiot savant” like



OpenCog: our unfair platform advantage

Given the current state of software and hardware tech, creating a human-level AGI is both

- A substantial research project, comprising multiple coupled sub-research-projects
- A large-scale software engineering project

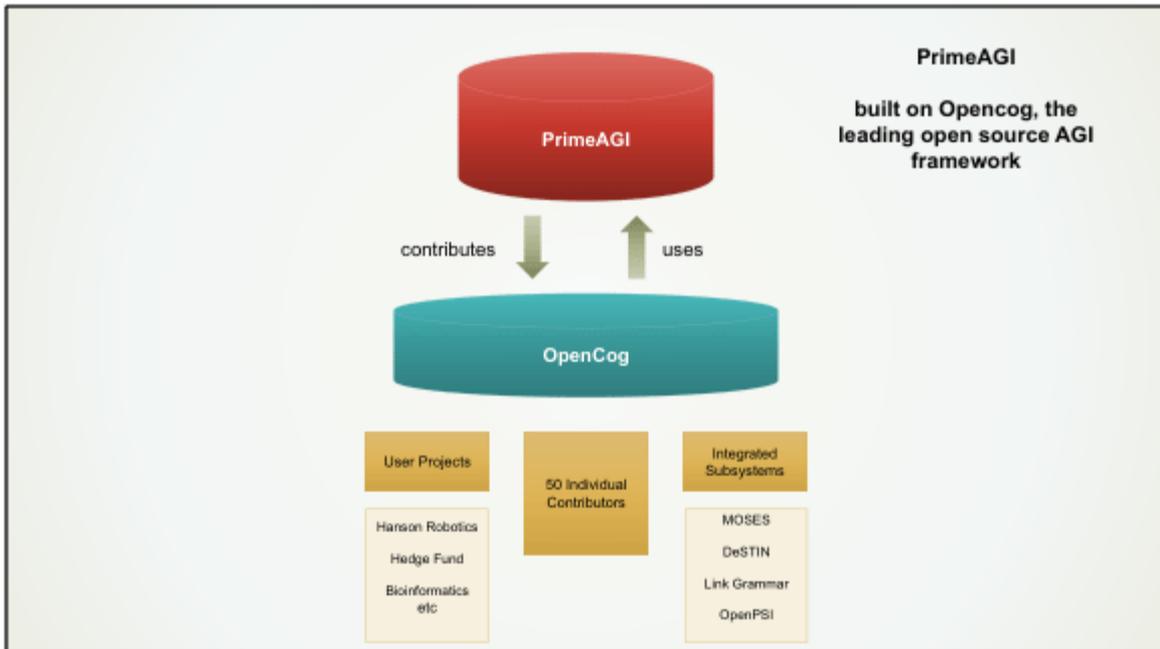
Executing on these two aspects in parallel presents unique challenges

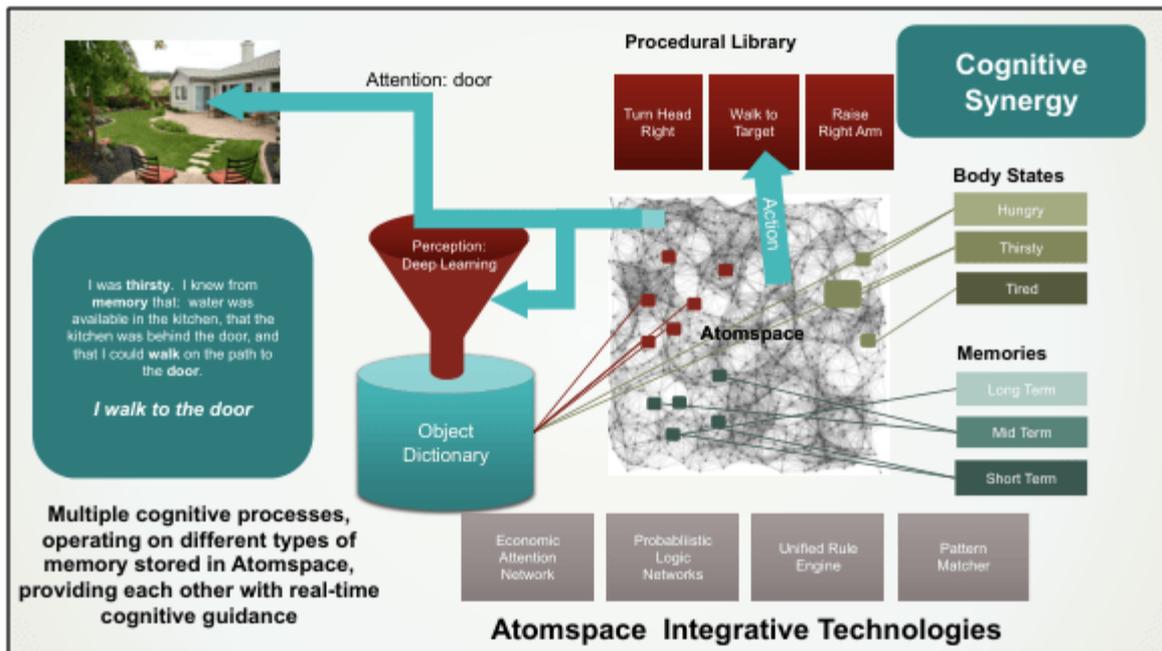
We address this issue via linked but separate development of

- **PrimeAGI**, a comprehensive, well fleshed out AGI design with deep theoretical grounding
- **OpenCog**, a software framework and platform designed for scalable implementation of hybrid AI systems – very suitable for PrimeAGI, but not only for PrimeAGI

Other proto-AGI projects lack a platform of comparable appropriateness and quality

- Academic projects are mainly implemented by grad students, with minimal attention to scalability, code quality, software architecture, etc.
- Corporate projects tend to be implemented on powerful distributed frameworks originally designed for non-AGI purposes, hence not fully apropos for the goal





1

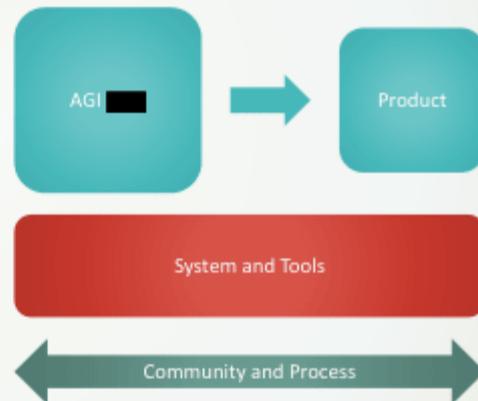
Resources and Timeline

- Aim: the "V2" PrimeAGI system; embodied dialogue with deep reasoning and understanding
- Approximately 20 full time technical staff, composed of four teams:
 - Artificial General Intelligence R&D
 - Product
 - System and Tools
 - Community and Process
- Ongoing collaborations with external researchers
- Three year project, with three phases, broadly described as setup, research, delivery.

1

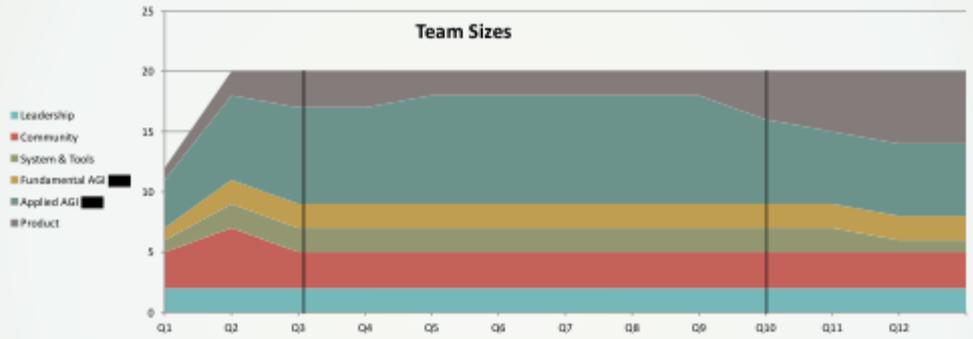
Team overview

- AGI R&D: Largest team, split into subteams (inference, language, learning, attention, cognitive control, perception, action, memory)
- Product: OSS releases, integration, tuning
- System & tools: infrastructure and architecture projects, visualization, etc.
- Community: demos, outreach, docs



1

Project Phases



1: Setup, 6 mos

- Hire and train
- Refine research specs
- OpenCog platform

2: Research, 21 mos

- Fundamental and applied research in each AGI topic
- Periodic integration and evaluation
- Continued improvements to platform, tooling and support

3: Delivery, 9 mos

- Final tuning
- Quantitative tests and evaluations
- Final product releases (Robot control, Minecraft)

1

Key Goals by Phase

1

- OpenCog ready to grow as an OSS organization
- Hire full team and bring them up to speed
- Refine research plans for fundamental research, and break experimental research projects into iterations

2

displaying actual understanding of the robot's interactions, context, people and of itself.
interactions used to drive behavior and guide learning so the robot is better able to please users and achieve goals
tion and inference modules driving goal achievement
that is able to demonstrate complex behavior and explain its actions and motivations
and maturity of the OpenCog platform and suite of OSS projects, attracting a wider community of AGI researchers and
ppers.

3

- Clear, strong demonstration of AGI achievements
- Secure funding for next major milestone
- Continue fundamental and applied R&D
- Continue OpenCog platform and OSS community growth

Core Team Members

- Dr. Ben Goertzel
- Jim Rutt, former Chairman of the Santa Fe Institute, CEO of Network Solutions, Inc., CTO Thomson-Reuters, investor and startup entrepreneur
- Cassio Pennachin, co-founder and chief architect of OpenCog, co-author of PrimeAGI design, Dr. Goertzel's AI collaborator since 1998
- Dr. Nil Geisweiller, PhD in probabilistic logic, developer on OpenCog and predecessors since 2005, co-author of PrimeAGI design
- Dr. Linas Vepstas, PhD in physics, developer on OpenCog and predecessors since 2007, lead OpenCog developer since 2006, lead developer for GnuCash, experience developing for Linux kernel and numerous major corporate and OSS projects
- Andre' Luiz de Senna, co-founder of OpenCog and lead developer of original OpenCog code, Goertzel and Pennachin's AI collaborator since 1999
- Dr. Eddie Monroe, PhD in psychology, developer on OpenCog and predecessors since 2006

Additional Material

- **To be added to the presentation as appendices, covering the following topics:**
 - Atomspace architecture and application
 - PrimeAGI design and concepts
 - Project plan specifics
 - Comparison to competing projects