

OpenCog AGI Toddler Project

**Open-Source “Human Toddler Level” AGI by 2015
for US \$3M Total**

Ben Goertzel and Gino Yu, Principal Investigators
Multimedia Innovation Centre, P507 Hong Kong Polytechnic University
Hung Hom, Hong Kong

Rough draft, not yet for distribution

Introduction	2
AGI: The Time is Now	2
Project Strategy	7
Four Year Plan	11
Organizational Structure and Management	15
Preliminary Budget	18

Introduction

The goal of the proposed “OpenCog Toddler” project is to achieve a radical breakthrough in Artificial General Intelligence, rapidly and at minimal cost. Specifically, the goal is to use the OpenCog design and codebase to create an AGI system with the rough general intelligence of a human 3-4 year old child, demonstrated via embodiment in virtual world characters and humanoid robots, within the scope of a 4 year R&D project.

The project will involve roughly 15 full time staff, which is achievable within US\$750K/year (including hardware, rent and other costs) -- a bargain price that is possible due to relatively low costs in the proposed location in Hong Kong Polytechnic University’s M-Lab technology incubator. The proposers are currently leading a smaller, related project in the M-Lab, the “OpenCog Game AI project”, with 6 programmers working on an OpenCog-based toolkit for the creation of intelligent video game characters. The OpenCog Toddler project team would be seeded with individuals already on location in the M-Lab and experienced on the OpenCog Game AI project, in addition to new hires from within Hong Kong and overseas. All software created would be open-source and made available to the world at large.

To support this initiative, we require US \$3M total, to be spent over 4 years, primarily on salaries.

While the software created in this effort will have massive commercialization potential, the proposed work specifically described here focuses on AGI R&D; development of commercial prototypes or products would require additional staff and funding.

AGI: The Time is Now

What do we mean by AGI, Artificial General Intelligence? And why do we believe the OpenCog Toddler project will have such a dramatic impact on the overall global progress of AGI R&D?

The proximal goal of the AGI field is simple, powerful and ambitious: **the development and demonstration of systems that exhibit the broad range of general intelligence found in humans**. In its most ambitious interpretation, “AGI” may be used to refer to an extreme generality of intelligence potentially going far beyond the human level; Marcus Hutter has modeled this mathematically with his theory of Universal AI. However, for the proposed project the focus will be on Human-Level AGI, and unless otherwise noted, that is what we will take the term “AGI” to mean here. The goal of developing AGI echoes that of the early years of the Artificial Intelligence movement, which after many valiant efforts largely settled for research into ‘narrow AI’ systems that could demonstrate or surpass human performance in a specific sort of task, but could not generalize this capability to other types of tasks or other domains.

A classic example of this “narrow AI” approach was IBM’s DeepBlue system, which could beat world chess champion Gary Kasparov but not apply that skill to any other problem domain (without substantial human re-programming). In 2010, IBM’s Watson question answering system dramatically defeated two all time champions in the quiz show *Jeopardy!*, but having never visited Chicago’s O’Hare and Midway airports, fumbled on a question that any human frequent flier would have known. And to apply the technology underlying Watson to another domain, such as medicine or call center support, would require not merely education of the AI system, but significant reprogramming -- the analogue of needing to perform brain surgery on a human each time they need to confront a new sort of task. As impressive as these and other AI systems are in their restricted roles, they all lack the basic cognitive capabilities and common sense of a typical three year old child, let alone a fully educated adult professional.

Nils Nilsson, one of the early leaders of the AI field, stated the goal of AGI quite crisply in the 2005 AI Magazine article “Human-Level Artificial Intelligence? Be Serious!”):

“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.”

Or, to get a more detailed feel for the broad scope of AGI, see Table 1, which was produced by a team of 12 AGI researchers at the *AGI Roadmap Workshop* organized by Ben Goertzel and Itamar Arel at the University of Tennessee, Knoxville in 2009. This table lists a number of the key competencies that an AGI system should be expected to display if it’s going to be considered a proper human-level AGI. By the end of the proposed 4 year OpenCog Toddler project, the OpenCog AGI system will be able to handle every one of these capabilities, at least in a simplistic way.

Early AI researchers, though their hearts were often in the right place, lacked sufficiently powerful computers and software tools and sufficient knowledge of cognition to practically succeed at their goals. Right now, in 2011, Artificial General Intelligence is an R&D domain whose time has come -- after a few decades of “AGI winter” it is finally beginning to flourish, with a host of research ideas, software systems and practical applications emerging from the minds of AGI researchers worldwide.

Table 1. Some of the Important Competency Areas Associated with Human-Level General Intelligence

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		

Creating the Sputnik of AGI

The success of the OpenCog Toddler project will have implications far and wide. At this stage, all that is needed to seed a massive global acceleration in the AGI field is one truly compelling demonstration of machine intelligence -- one working, vividly demonstrable software system that both embodies a serious scientific approach to AGI, and also cries out to the naive viewer "This has the feel of being on the path to human-level general intelligence!!!" Such a demonstration would serve as a "**Sputnik of AGI**", inspiring others to put their energetic and financial resources into pursuing AGI as well.

One of our goals with the proposed project is to create precisely this "Sputnik of AGI"

Specifically, what aspect of the OpenCog Toddler will have this Sputnik-like impact? The main point is not to demonstrate any particular functionality, but rather to **demonstrate the capability to robustly learn new behaviors** (where this learning occurs via experimenting with the world and via communicating with human beings, not via carefully supervised training on human-prepared training sets).

We aim to demonstrate this learning capability in the context of virtual world characters and humanoid robots. (By a "virtual world" here we mean a 3D video game type world, but without necessarily having the game-play associated with a video game.) The dual focus on virtual characters and humanoid robots is necessary because of the limitations of each of these modalities. Virtual worlds lack the richness of the physical world, yet current robots lack robust sensors and manipulators. By using the same OpenCog AGI mind to control both sorts of agents, we can overcome these limitations; the OpenCog system is capable of integrating what it learns via its robot body with what it learns via its virtual embodiment, into a single unified "mind-space."

The Big Picture: Beyond the AGI Toddler

While an "AGI toddler" will be a massive R&D achievement, obviously it is not of particular practical value on its own. The OpenCog project has developed a longer-term roadmap, in which the proposed 4 years of effort is merely the next phase. We are focused on this phase currently, both because it's the next natural step in our R&D work, and because we believe that it can serve as a "Sputnik of AGI" -- so that, once a robust AGI toddler is demonstrated, large portions of the world will wake up to the viability of achieving more and more powerful AGI in the near term. We strongly suspect that, with a robust AGI toddler in hand, a large array of commercial and government sources will find it in their interest to pursue further AGI development toward a large variety of practical and research goals.

The next step beyond the AGI Toddler will be to further scale up and educate the OpenCog AGI system, enabling it to cope with more complex learning such as the typical elementary school curriculum, and with more complex physical tasks such as

would be required of a household service robot. This next step will enable the development of a variety of “AGI specialists” carrying out tasks of humanitarian and commercial value -- for instance, service robots, automated scientists, automated call center agents, generally intelligent financial trading systems, and so forth (the list is almost endless). Integration of these specialists within the integrative OpenCog architecture will ultimately lead to a generally intelligent software system with the common sense of an ordinary human, but a host of specialized capabilities going beyond the human level. On the overall OpenCog roadmap, we conjecture that this stage of “Human Level Plus” AGI may be achievable around 2023. Beyond this level, further development will also be likely, but becomes even more difficult to chart in details -- one gets into the territory of what Vernor Vinge, Ray Kurzweil and others have called “The Singularity.”

While this long-term picture is much of our motivation for pursuing the proposed R&D, nevertheless in the context of the proposed 4 year project, our focus will be squarely on achieving an AGI Toddler using OpenCog. Even starting from a well-articulated design and a well-built codebase, as we are, this a very difficult technical challenge. Once we have met this challenge, we will proceed to the next steps on the broader OpenCog roadmap, leading toward ever more powerful AGI.

Use of Funds

To ensure that the proposed work is carried out effectively, the resources that we request will be applied towards the following:

- **Management and Technical Leadership** - The largest portion of the funding will be used to pay experienced AI researchers and developers to carry out the proposed R&D program. Technical leadership will be provided by a group of individuals already experienced with the OpenCog project
 - *Dr. Ben Goertzel, Executive Director*, primary author of the OpenCogPrime AGI design and chief OpenCog project founder
 - *Dr. Joel Pitt, Director of AGI Development*, OpenCog project co-founder and team leader of current OpenCog Game AI project in Hong Kong
 - *Dr. Nil Geisweiller, Chief Scientist*, with 4 years of OpenCog experience
 - *Dr. Zhenhua Cai, Lead AGI Developer*, senior AI developer in current OpenCog Game AI project [to receive PhD in 2012]
 - *Dr. Ruiting Lian, Lead NLP Developer*, lead developer of OpenCog’s language generation and dialogue systems, to join OpenCog Game AI project in Fall 2012 [to receive PhD in 2012]
 - *Dan Zwell, Director of Environments & Evaluation*, experienced commercial software developer, previously worked on video game AI using the Novamente Cognition Engine, a direct predecessor of OpenCog
- **Technical Staff** - The above technical leaders will be assisted in their work with 6 AGI software developers focused on OpenCog development (two of whom, Jared Wigmore and Deheng Huang, are current members of the OpenCog Game AI team), and 2 developers focused on the virtual world and robotics aspects.

- **Facilities** -- We anticipate physically locating the project in the M-Lab, a tech incubator associated with Hong Kong Polytechnic University, where the OpenCog Game AI project is currently situated
- **Computing Hardware** – The OpenCog Toddler project will require a modest amount of advanced computer hardware, including development machines for staff, a network of ~5 powerful multiprocessor servers, and a GPU supercomputer such as an Nvidia Kepler. Nevertheless, given the rapidly decreasing cost of advanced computing hardware, this is a relatively small part of the total budget, estimated at < 10%.
- **Robotics Hardware** -- In the OpenCog lab at Xiamen University, we have worked with the Aldebaran Nao robot; but for the OpenCog Toddler project we are more excited about the new Hanson Robotics RoboKind. These are small humanoid robots with high-quality cameras, relatively capable hands, and humanlike facial expressions, available for \$15K each. We will begin by purchasing two of these for the project (two so as to enable experimentation with social interaction), and then purchase further robots for the project in subsequent years as robotics technology advances.

Project Strategy

The foundations of the proposed OpenCog Toddler project are

- **the OpenCogPrime design for machine intelligence**, spelled out in detail in the 900-page multi-author manuscript *Building Better Minds* (which currently exists as a rough draft available to interested parties, and is being edited and refined with a view toward publication in 2012).
- **the OpenCog codebase**, which currently implements roughly 40% of the overall OpenCog design (see <http://opencog.org>). It has been under development since 2008, building on code provided by Novamente LLC that was under development since 2001. Most of the code is C++, but recently a significant percentage of development being done in Python on top of the C++ core

In the proposed project, the remaining 60% of the OpenCogPrime design will be implemented in the OpenCog codebase, and used for the control of virtual-world and robotic agents, with a focus on enabling robust learning via experience and communication.

OpenCog

OpenCogPrime (OCP) is a comprehensive architecture for cognition, language, and virtual agent control, created by a team of AI experts led by Dr. Ben Goertzel during the period since 2001 (and building on their work from the 1990s). Conceptually founded on the systems theory of intelligence summarized in the 2006 book *The Hidden Pattern* (and in more detail in the forthcoming book *Building Better Minds*), it is currently under development within the open-source OpenCog AI framework. The architecture combines multiple AI paradigms such as uncertain logic, computational linguistics, evolutionary program learning and connectionist attention allocation in a unified

cognitive-science-based architecture. Cognitive processes embodying these different paradigms interoperate together on a common neural-symbolic knowledge store called the Atomspace.

OpenCog has been used for numerous commercial applications in the area of natural language processing and data mining. For instance, in work with the Clinical Center of the US National Institutes of Health, OpenCog's PLN reasoning and RelEx language processing were combined to do automated biological hypothesis generation based on information gathered from PubMed abstracts. Data mining applications have spanned text classification, bioinformatics and computational finance. OpenCog has also been used to control virtual agents in videogame type worlds (including current work in the M-Lab, and earlier work controlling virtual dogs in Second Life and Multiverse) and humanoid robots (in the BLISS Lab at Xiamen University).

While OpenCog is not closely based on the human brain, it derives its overall architecture from cognitive science, and as such it shares some general properties with human intelligence. One of these properties is a fundamental complexity that makes the design difficult to summarize in one or two sentences. The brain consists of a host of different regions, operating on similar principles yet with significant differences also, each carrying out separate but overlapping functions, and interoperating with each other constantly in subtle ways. Very few neuroscientists believe there is a single, simple algorithm or structure underlying the brain's intelligence -- rather, there seems to be a variety of interacting, complexly interrelated mechanisms. Similarly, OpenCog is a complex design, which contains multiple memory stores corresponding roughly to the types of human memory identified in cognitive science, and learning algorithms corresponding to each type of memory.

One of the key principles of the OpenCog design is cognitive synergy, a technical concept meaning very roughly that the different learning algorithms must interact in such a way as to help each other scale to large problems. On its own, any one of OpenCog's learning algorithms would choke on large problems, due to what computer scientists call "combinatorial explosion" (the existence of far too many combinations of observations and data items and conjectures, for any algorithm to be able to blindly search through them all). However, the overall architecture is designed in such a way that, when one algorithm begins to run into a combinatorial explosion, the others can help it out. It seems likely that the human brain obeys a similar conceptual principle, though manifested in different ways. The various memory stores and algorithms in OpenCog all use a probability-theoretic semantics to represent knowledge (sometimes among other means), which provides a clean and simple framework within which cognitive-synergetic dynamics may operate. Tables 2 and 3 below briefly review some of the key knowledge structures and cognitive algorithms in the system.

Table 2. Key Memory Types and Corresponding OpenCog Knowledge Structures

Memory Type	OpenCogPrime Knowledge Structure
Declarative	The AtomTable, which is a special form of weighted, labeled hypergraph -- i.e. a table of nodes and links (collectively referred to as Atoms) with different types, and each weighted with a multi-dimensional truth value (embodying an "indefinite probability" value that give both probability and confidence information).
Attentional	Atoms in the AtomTable are weighted with AttentionValue objects, which contain both ShortTermImportance values (governing processor time allocation) and LongTerm Importance values (governing memory usage).
Procedural	This is handled using special "Combo" tree structures embodying LISP-like programs, in a special program dialect intended to manage behaviors in a virtual world and actions in the AtomTable
Sensory	Handled via a collection of specialized sense-modality-specific data structures
Episodic	Handled via an internal simulation world that allows the system to run "mind's eye movies" of situations it remembers, has heard about, or hypothetically envisions.
Intentional	Goals are represented by Atoms stored in the AtomTable; there is a separate table indicating which Atoms are top-level goals, which is used to guide attention allocation and goal refinement processes

Table 3. OpenCog's Key Cognitive Algorithms

Cognitive Process	OpenCogPrime algorithm
Uncertain inference	Probabilistic Logic Networks (PLN) , a logical inference framework capable of uncertain reasoning about abstract knowledge, everyday commonsense knowledge, and low-level perceptual and motor knowledge
Supervised procedure learning	MOSES , a probabilistic evolutionary learning algorithm, which learns procedures (represented as LISP-like program trees) based on specifications
Attention allocation	Economic Attention Networks (ECAN) , a framework for allocating (memory and processor) attention among items of knowledge and cognitive processes, utilizing a synthesis of ideas from neural networks and artificial economics. ECAN also comes with a forgetting agent that either saves to disk or deletes knowledge that is estimated not sufficiently valuable to keep in memory.
Map formation	Use of frequent subgraph mining, MOSES and other algorithms to scan the knowledge base of the system for patterns and then embodying these patterns explicitly as new knowledge items
Concept creation	A collection of heuristics for forming new concepts via combining existing ones, including conceptual blending, mutation and extensional and intensional logical operators
Simulation	The running of simulations of (remembered or imagined) external-world scenarios in an internal world-simulation engine
Goal refinement	Transformation of given goals into sets of subgoals, using concept creation, inference and procedure learning
Perception and Actuation	Sensory and motor processing are handled in OpenCog via the DeSTIN architecture, a software system originally created by Dr. Itamar Arel (at the University of Tennessee, Knoxville), which uses hierarchical spatiotemporal memory networks to enable scalable perception, state inference and reinforcement-learning-guided action.

Environments and Evaluation for Generally Intelligent Agents

One of the most important issues in any long-term research project is how to gauge incremental progress. How do we know we're getting closer and closer to our end goal of advanced AGI?

The commonplace practice of identifying highly specific tasks and then comparing the performance of various systems on these tasks, doesn't work well for AGI, because for any highly specific task there is likely to be a reasonably effective "narrow AI" solution. The crux of AGI is the capability to generalize and perform broadly beyond specific tasks; but this is a more difficult thing to test.

These issues were discussed in depth at the 2009 AGI Roadmap Workshop, mentioned above. Among the key conclusions reached there was that, even though it's difficult to formulate precise tests that are naturally applicable across different AGI paradigms, it's less problematic to identify *environments* and *tasks* that multiple AGI researchers can agree are appropriate.

With this in mind, our approach in the proposed project is to create environments, with associated task sets, suitable for evaluating the behavior of various approaches to AGI, and assessing their progress toward advanced AGI. The goal is not to posit any single "AGI IQ test" but rather to create tools suitable for evaluating AGI systems according to a variety of appropriate criteria. Refining the details of the task sets is something the Scientific Advisory Board will be very helpful with.

Videogame Agents

Our current OpenCog Game AI project at the M-Lab involves using OpenCog to control a game character in a game world loosely modeled on the commercial game Minecraft (in which a wide variety of objects and scenarios may be created via manipulating small blocks). The game design and engineering aspect of this work will be extended to create a more general Minecraft-like environment, suitable for testing the full variety of cognitive capabilities expected of a human-level AGI. Nearly the full spectrum of intelligence testing tasks used to evaluate the cognitive capabilities of young human children may be implemented in this context. In addition to its research value, this evaluation environment possesses obvious value in terms of assessing the readiness of AGI systems for deployment in a commercial gaming context.

Humanoid Robotics

The OpenCog system is now being used to control a Nao humanoid robot, in the Brain-Like Intelligent Systems (BLISS) Lab at Xiamen University, in China. The interface between OpenCog and the Nao uses the ROS Robot Operating System created by

Willow Garage, which may easily be utilized for other AI systems besides OpenCog, and for other robots besides the Nao. Via refining and extending the software currently used to connect OpenCog to the Nao, we can relatively easily create a general framework for AGI-robot interaction.

Given a general interface of this nature, what remains to create a robust environment for evaluating AGI systems in a robotics context, is to create a robot lab environment providing sufficiently rich interactions within the restrictions posed by the sensors and actuators of affordably available robots (such as the Hanson Robotics Robokind). This is a relatively simple matter, requiring only some systematic experimentation regarding the available robots' precise capabilities. Such an environment would have clear utility for research and also for the evaluation of the readiness of AGI systems to control toy or household service robots in a commercial setting.

Four Year Plan

A high-level plan for the proposed four-year project is as follows, assuming project initiation in January 2012.

Year 1: 2012

	Key Accomplishments in Year 1
Example internal OpenCog achievements	<ul style="list-style-type: none"> • Simple English dialogue regarding scenes in the virtual world • Complete integration with DeSTIN for vision processing, enabling basic object and event recognition in the robot lab • Complete integration with ROS for flexible robot control in robot lab environment, allowing execution of complex motor plans • Tune inference engine for spatiotemporal inference in videogame and robot lab contexts, allowing commonsense reasoning about simple situations
Environments & Evaluation	<ul style="list-style-type: none"> • Completion of video-game world environment for evaluating AGI systems • Implementation of simple (ROS based) humanoid-robot API enabling AGI systems to easily interact with robots • Outfitting of simple humanoid robot lab for AGI experimentation
Example behavioral milestone achievements	<ul style="list-style-type: none"> • Ability to hold simple English conversations about what its virtual agent or robotic environment sees and does, evincing basic "commonsense understanding" in the limited contexts of its virtual world and robot lab • The scope of conversations will include all the major types of speech act delineated by Speech Act theory • In the virtual world: ability to build structures from blocks to help it achieve its goals, e.g. hiding from an attacker, reaching a high-up object, building a bridge between one place and another

Year 2: 2013

	Key Accomplishments in Year 2
Example internal OpenCog achievements	<ul style="list-style-type: none">• Completion of an integrated, OpenCog-based “proto-AGI mind”• Ability to improve language understanding and generation capability via experience• Robust question answering regarding situations in the virtual world and robot lab, based on integration of inference with other cognitive functions• Formation of new concepts in response to new situations and environments• Integration of knowledge from large databases with experiential knowledge
Environments & Evaluation	<ul style="list-style-type: none">• Implementation of a spectrum of intelligence-testing tasks within the video-game environment• Completion of the creation of a robot-lab environment suitable for AGI experimentation
Example behavioral milestone achievements	<ul style="list-style-type: none">• When the agent is put in new situations, it is able to form new concepts as needed for achieving simple goals, and then use simple English to generate questions or answers about these concepts• Example: If the agent has never seen a small animal before (only humans, robots and inanimate objects), then after watching it for a while, it may be asked to follow, attract or capture the animal, and describe what it's doing. (Obviously the specifics of this example would be different in the virtual world and robot lab cases.) If it's then exposed to a second kind of small animal, it should be able to appropriately generalize what it has learned from interacting with the first one.• Ability to build blocks structures in the virtual world that imitate things it has seen in pictures, or has been shown in the virtual world• Ability to build simple blocks structures in the robot lab

Year 3: 2014

	Key Accomplishments in Year 3
<p>Example internal OpenCog achievements</p>	<ul style="list-style-type: none"> • Increase in the system's learning and reasoning capability, leading to more useful performance on practical tasks • Use of abstract inference to draw new conclusions combining knowledge derived from experience with knowledge derived from databases and text corpuses • Scalable implementation on a reasonably large distributed network • Feedback enabling cognition to guide the details of perception and action • Generalization of knowledge and intuition back and forth between virtual-world and robot domain
<p>Environments & Evaluation</p>	<ul style="list-style-type: none"> • Systematic intelligence testing of AGI systems in the video-game and data-analysis environments • Implementation of a suite of intelligence-testing tasks in the robot lab environment
<p>Example behavioral milestone achievements</p>	<ul style="list-style-type: none"> • If the agent is shown an object it has never seen but has only been told about verbally, or read about, it can use this knowledge to figure out what to do with it. • If the agent is shown an object in the virtual world and learns what to do with it, then when it's shown a similar object in the robot lab, it can quickly understand it ... and vice versa • In the virtual world: Ability to build complex blocks structures collaboratively with others, to achieve goals • Ability to build simple blocks structures in the robot lab • Ability to describe complex blocks structures and other physical or virtual entities and situations using complex sentences with multiple phrases and clauses, and series of linked sentences. • In the virtual world: Ability to learn new, reasonably complex behaviors via following instructions, asking questions when needed

Year 4: 2015

	Key Accomplishments in Year 4
Example internal OpenCog achievements	<ul style="list-style-type: none"> • Robot control outside the “robot playroom” in a broader indoor office context • Creative invention in the virtual world AI context • Flexible natural language dialogue, involving understanding and production of complex sentences
Environments & Evaluation	<ul style="list-style-type: none"> • Systematic intelligence testing of AGI systems in the video-game and data-analysis and robot lab environments • Addition of new intelligence tests to the various environments, in order to test the more advanced levels of intelligence that the AGI systems have achieved
Example behavioral milestone achievements	<ul style="list-style-type: none"> • In the robot lab: ability to build moderately complex blocks structures and describe and answer questions about them (this is basically a matter of having sufficiently advanced vision and motor control, as the cognitive aspects would have been achieved in the virtual world long before) • Ability to construct new entities using blocks and other available materials, qualitatively different from things it’s seen before, and discuss what it has done and why • Ability to encounter new entities, ask questions about them, and figure out how to practically utilize them based on the answers • Ability to learn from a combination of observation and dialogue, how to carry out new activities like (e.g.) tag or hide-and seek, or particular sorts of collective blocks play

In reflecting on the behavioral milestones in the above tables, it’s important to recall that the project goal is *not* to specifically tune the system for a set of specific milestone tasks each year. This would produce an interesting virtual world or robot toddler system, but not a platform for the creation of increasingly powerful AGI. Rather, the goal is to create a progressively more intelligent proto-AGI system using the OpenCog design and codebase, and use it to control a virtual agent and humanoid robot, and then test and demonstrate its capabilities using an assemblage of demonstrable and measurable tasks, like the ones outlined above. The real point is not the specific tasks but rather the cognitive capabilities underlying them, which will constitute the same sort of basic “commonsense understanding” that underlies a human toddler’s thinking -- and will then serve as a platform for further AGI development, in a multitude of scientifically, pragmatically and commercially valuable directions.

Organizational Structure and Management

Business Structure

The possibilities exist to structure the OpenCog Toddler via either a non-profit or for-profit enterprise.

On the nonprofit side, it would be straightforward to affiliate the project with a university (e.g. Hong Kong Polytechnic University, that hosts the current OpenCog Hong Kong project focused on video game AGI); however, in practical terms this affiliation would result in extra expenditure due to university overheads. The M-Lab tech incubator, where the current OpenCog Hong Kong project is focused, hosts both university and non-university projects. So the most efficient nonprofit route would likely be to host the proposed project in the M-Lab, but funded via an external nonprofit rather than by a grant to the university. Two nonprofits that could potentially serve this role would be the OpenCog Foundation, or Humanity+ (a US nonprofit with 501(c)-3 status, which operates in close collaboration with the OpenCog Foundation). Project leader Dr. Ben Goertzel holds Board-level roles in both these organizations. In any case, the foundation selected would then transfer funds to the M-Lab periodically (e.g. quarterly) as needed.

On the for-profit side, it would be simple to operate the project via Novamente LLC, an AI consulting and research company founded and operated by proposer Dr. Ben Goertzel. The OpenCog codebase is an outgrowth of code initially created by Novamente LLC during 2001-2008 under the name "Novamente Cognition Engine," and then open-sourced in 2008; since that time Novamente LLC staff have continued to contribute to the OpenCog codebase, and Novamente is a co-sponsor of the OpenCog Hong Kong Game AI project.

Potentially, it would also be viable to create a new firm to operate the OpenCog Toddler project, with a variety of business models, including game AI, toy robotics, service robotics, or AGI more broadly. It would also be viable to operate the proposed project via some other existing commercial entity. We are committed to preserving the open-source nature of the project, as we strongly feel this is best for promoting AGI progress overall. However, a variety of different business models and formal structures are consistent with OSS software development, and we are open to exploring these.

Organizational Structure

The formal organization structure suggested for the proposed project is outlined below. The upper levels of the org chart may vary slightly depending on the business structure, but almost everything from the Executive Director down should be independent of business structure.

- Management Board
- Scientific Advisory Board
- Executive Director: *Dr. Ben Goertzel*
 - Chief Scientist: *Dr. Nil Geisweiller*
 - Director of AGI Development: *Dr. Joel Pitt*
 - Lead AGI Developer: *Dr. Zhenhua Cai*
 - Lead NLP Developer: *Dr. Ruiting Lian*
 - AGI Developer: *Jared Wigmore*
 - AGI Developer: *Deheng Huang*
 - 4 additional AGI Developers: *TBD*
 - Director of Environments and Evaluation: *Daniel Zwell*
 - Virtual-World Software Developer: *TBD*
 - Robotics Software Developer: *TBD*
 - Software Tester: *TBD*
 - Computer System Administrator: *TBD*
 - Administrative Assistant: *TBD*

Management Board

The management board will meet biannually to monitor the progress of the project, and will approve all new Director-level staff hired under the initiative's budget. The management board will initially be composed of Dr. Ben Goertzel, Dr. Gino Yu and one or more representatives appointed by the project donors.

Scientific Advisory Board

The project's Scientific Advisory Board (SAB) will world's leading AGI researchers, selected due to the overlap of their research with the specific R&D being pursued at the in the project. A small amount of funds will be set aside each year to allow SAB members to visit and collaborate on research

The initial members of the Scientific Advisory Board are:

- Itamar Arel (University of Tennessee, Knoxville)
- Joscha Bach (University of Humboldt)
- Pei Wang (Temple University)
- Paul Rosenbloom (University of Southern California)

with a handful of additions expected.

Executive Director

The project's Executive Director, Dr. Ben Goertzel, will lead the overall project, and also participate in research and development as appropriate. He will also assist the

Directors of AGI Development and Environments & Evaluation in project management of their teams.

Chief Scientist

The central role of the Chief Scientist, Dr. Nil Geisweiller, will be to ensure that the AGI development occurring is done in accordance with the OpenCog plans and documentation; and to appropriately update the OpenCog plans, theory and documentation in accordance with what is learned in the course of the project. These responsibilities will not generally be a full-time occupation and hence Dr. Geisweiller will also participate in hands-on AGI development.

Director of AGI Development

The Director of AGI Development, Dr. Joel Pitt, will lead the implementation work on the project, including the allocation of tasks to the other AGI developers and the review of code and validation of successful task completion. He will also participate in AGI software design and development as time permits!

AGI Software Developers

The heart of the project team will consist of 8 AGI software developers working together to implement the OpenCog AGI design for intelligent agent control. Half of these developers are already experienced with OpenCog development, including Lead AGI Developer Zhenhua Cai, and Lead NLP Developer Ruiting Lian (both of whom will complete their PhDs in 2012).

Director of Environments & Evaluation; Virtual Worlds Developer; and Robotics Developer

Creating environments for AGI agents to live and learn in, and tests to evaluate AGI agent intelligence, are significant tasks independent of AGI development itself. The Director of Environments & Evaluation will coordinate this effort, playing a technical lead and project management role, assisted by two junior software developers, one focusing on the virtual world environment and one on the robots and robot lab. This role will initially be played by Daniel Zwell, who in 2007 worked on the predecessor to OpenCog within Novamente LLC, and has more recently been leading a small commercial software development team in Shanghai.

Software Testing and System Administration

So that the AGI software developers may focus strictly on AGI, the team will be supported by a full-time software tester (who will perform a combination of bug testing and intelligence testing) and a system administrator (who will assist with maintaining the robots as well as the computer hardware, and also deal with any issues regarding maintenance of the code and documentation repositories).

Preliminary Budget

A preliminary budget for the proposed project is as follows:

Category	Annual Cost in USD
Salary (incl. pension, medical)	\$575K
Computing hardware	\$50K
Travel for Scientific Advisory Board visits	\$15K
Travel for Project Staff	\$10K
Office (rent, AC, etc.)	\$70K
Robots	\$20K
Accounting & Legal	\$10K
Total	\$750K

with a detailed breakout of the main expense (salaries) below. Specific salary numbers for named individuals are withheld below for privacy reasons, but have been estimated somewhat carefully. Further financial details may be shared with serious potential donors or investors upon request.

Title	Name	Estimated annual salary in USD before overheads
Executive Director	Dr. Ben Goertzel	***
Chief Scientist	Dr. Nil Geisweiller	***
Director of AGI Development	Dr. Joel Pitt	***
Director of Environments & Evaluation	Daniel Zwell	***
Lead AGI Developer	Dr. Zhenhua Cai	***
Lead NLP Developer	Dr. Ruiting Lian	***
AGI Developer	Jared Wigmore	***

Title	Name	Estimated annual salary in USD before overheads
AGI Developer	Deheng Huang	***
AGI Developer	TBD	\$25K
Virtual World Software Developer	TBD	\$25K
Robotics Software Developer	TBD	\$25K
Software Tester	TBD	\$15K
System Administrator	TBD	\$15K
Administrative Assistant	TBD	\$10K
Total Salary Before Overhead	--	\$520K
Total Salary Including Pension & Health Insurance	--	\$575K

Computing hardware needs will be determined as the project unfolds, but are roughly envisioned to include:

- Development laptops for each developer that needs one (to be replaced in 2014)
- At least 2 multiprocessor servers with generous RAM purchased per year (e.g. a current OpenCog server has 16 processors and 96BG RAM)
- In 2012, an Nvidia Kepler GPU supercomputer; and potentially updated versions later as they are released

Regarding robots, in 2012 we will purchase 2 Robokind from Hanson Robotics, at a cost of roughly \$15K each. Further robotics purchases will depend on the technology available, which is hard to foresee in advance, as robotics is a rapidly developing field. Possibilities include augmenting the Robokind with additional sensors such as artificial skin as this becomes available.