Creating Robots with Toddler-Level Intelligence Using the OpenCog AGI Architecture

Research Proposal for the Epstein Foundation

Ben Goertzel
February, 2013

Introduction

At its inception in the 1950s, the AI field aimed at producing human level general intelligence in computers and robots. Within a decade or so the difficulty of that goal became evident, and the AI field refocused on producing systems displaying intelligence within narrow domains. This focus on "narrow AI" has been strikingly successful in some regards, leading to practical AI applications such as Google's search and ad engines, Deep Blue and other game-playing AIs, IBM's Watson Jeopardy-player, a host of profitable AI financial trading systems, and so forth. Over the past few years, however, there has been a resurgence of research interest in the original goals of AI, often using terminology such as Human-Level AI or Artificial General Intelligence (AGI) [1,2,3,4]. The core reason for this resurgence is a feeling that, due to advances in the AI field and in allied areas such as computer and robotic hardware, computer science, cognitive psychology and neuroscience, we are in a far better position to approach these goals today than were the founders of AI in the 1950s.

One may ask why, given all the amazing recent developments in applied AI and allied areas, we have not yet seen AI software systems with humanlike general intelligence. We believe there is one key ingredient missing – the effective linkage of **subsymbolic** AI methods, dealing with raw perceptual and motoric data, with **symbolic** AI methods, dealing with abstract reasoning and language, and higher level cognition. The AI field now possesses able algorithms and architectures on both the symbolic and subsymbolic sides, but without both aspects working together, human level general intelligence is hard to come by.

Some researchers aim to bridge the gap by making subsymbolic AI systems intelligent enough that they can learn symbolic reasoning via experience. After all, they figure, symbolic reasoning originally evolved from subsymbolic thinking – humanity's distant evolutionary ancestors probably didn't do much symbolic reasoning. Others aim to bypass the need for subsymbolic processing, figuring that one can create a human-level AI well enough by communicating with it using text chat, and having it gather knowledge from the Web and structured databases – all sources that are easier to feed directly into a symbolic AI system. Some feel the crux of the AI problem lies on the symbolic side, and if one wants to make an AI system controlling a robot, one can simply bolt some separate perception and motor-control modules onto one's symbolic AI system.

The solution we suggest is different, and in a certain sense simpler. We propose to interconnect a highly functional, primarily symbolic AGI architecture (OpenCog [5], an international open source project[1]), with a highly functional subsymbolic AI system (DeSTIN [6,7], developed at the University of Tennessee, Knoxville). We propose to perform this connection, not merely by linking the two systems as separate modules, but by enabling the two systems to exchange deep information regarding their internal states, and provide guidance to each other's thinking. In order to do this, we have designed a unique pattern recognition layer intended to live between DeSTIN and OpenCog, and translate between the languages of the two different systems [8]. The combined OpenCog/DeSTIN system will powerfully display a core principle of AGI called "cognitive synergy," key to the OpenCog architecture, according to which different aspects of an intelligent system are engineered to help each other out of cognitive bottlenecks.

To refine and test our approach to OpenCog/DeSTIN integration, we will pose the integrated system the task of controlling a Hanson Robokind robot in a robot lab environment – a cognitively enabled robot or "CogBot." We will specifically aim at enabling the Robokind to carry out a variety of preschool-like behaviors, such as playing with blocks, interpreting pictures, drawing with a pencil or marker, answering questions and following instructions.

This forms a natural extension of our current work using OpenCog to control animated characters that build things with blocks in a 3D video game world; and of prior work using OpenCog to control a Nao robot in a robot lab in a simpler way, without benefit of hybridization with DeSTIN or any similarly sophisticated perception/action system [9]. The use of a preschool-like setting enables the application of ideas from developmental psychology to guide and assess the AGI's progress [10].

[1] http://opencog.org

We don't aim at this stage to create a robot giving a perfect simulation of a human child. Both the body and mind of our proposed "CogBot" system are very different from those of any human being. Qualitatively, our aim with the robotics aspect of the project is to create a robot that is recognizably and irrefutably generally intelligent, in the rough manner of a young human child.

In order to rigorously evaluate our progress, we will make use of the WPPSI (Wechler Preschool and Primary Scale of Intelligence) test, the generally accepted IQ test instrument for children aged 2 through 7. This test may be carried out via a robot sitting at a table across from a human examiner, answering questions, writing on pieces of paper, looking at pictures, and manipulating blocks and puzzle pieces. It requires vision and simple physical manipulation of objects on a tabletop. It does not require carrying of objects, nor navigation (though as it happens OpenCog is already quite good at navigation). It also does not seek to measure social or emotional intelligence, though there do exist evaluation instruments for this such as the EQI-YV (Emotional Quotient Index – Youth Version), which may be interesting to explore in future work.

Due to the complexity of the underlying AI systems (OpenCog and DeSTIN), the proposed project is a multidisciplinary effort involving component problems in multiple AI areas, including computer vision and audition, humanoid robot control, computational linguistics, probabilistic reasoning, automatic program learning, assignment of credit and concept creation. Achieving the project goals will not require breakthroughs in any of these areas; the focus will rather be on integration and synergetic behavior. However, the results of this research are expected to yield interesting advances in each of these areas, in addition to the advancement toward human-level artificial general intelligence implicit in the achievement of childlike intelligence in a humanoid robot.

Existing & Desired Additional Funding

Funding for the work required to complete this project is already partially in place, due to an ITF (Innovation in Technology Fund) grant obtained via the Hong Kong government, with Dr. Gino Yu as Principal Investigator and Ben Goertzel's AI consulting firm Novamente LLC as corporate sponsor (providing 10% of the funds to match the ITF's 90%). This ITF grant is titled "Artificial Intelligence Software Enabling Toy Robots to Learn, Communicate, Emotionally Bond and Display Individual Personalities" (ITS/178/12FP). The ITF grant proposal states that the software produced with the government funding will be open source, with the Hong Kong government or university not having intellectual property rights.

The existing government funding amounts to US$342K from the ITF and $38K from Novamente LLC. However, the government funding is not sufficient to enable the OpenCog controlled robot to be brought to the level required to perform well on all the WPPSI test tasks, in a genuine way not based on explicitly engineering or teaching "to the test."

This government funding will pay for a team of 6 PhD students and junior programmers working in a university lab; the requested additional funding would cover

- A full-time senior software developer
- A full time, experienced AI PhD, to carry out and help supervise AI software development
- A part-time senior systems administrator
- A small, dedicated compute cluster (of Linux machines)
- A dedicated office / robot lab space for the project
- Assistance of Novamente LLC with its corporate contribution (which would be valued due to Novamente's limited funds and limited commercial activity currently)

In short, the additional funding would transform the project from an underfunded university-research-lab type initiative, into a professional R&D initiative with experienced hands-on leadership, thus at least doubling productivity and significantly increasing odds of success.

Background

The background for the project comes from multiple disciplinary directions – computer science, cognitive science, robotics and systems theory – and the project's conceptual direction is the result of collaboration between three interdisciplinary researchers and technologists:

- **Dr. Ben Goertzel**: Principal Investigator; scientific leader of commercial project sponsor Novamente LLC. Dr. Goertzel is an AI researcher with a mathematics background, the leader of the international Artificial General Intelligence research community.

- **Dr. David Hanson**: Co-Investigator; scientific leader of commercial project sponsor Hanson Robotics. Dr. Hanson's whose interdisciplinary background in arts, design and robotics led him to develop a variety of revolutionary robots including Robot Einstein and Hanson Robokind (the latter to be used in the proposed project) [see Figure 2]. These robots feature novel robot skin enabling humanlike emotional expression, and feature a unique combination of hardware and software features focused on rich emotional interaction with human users.

- **Dr. Itamar Arel**: Scientific Advisor. AI and engineering professor at the University of Tennessee Knoxville, the creator of the DeSTIN (Deep SpatioTemporal Inference Network) software to be used in the project for robot vision and action (to serve as an intermediate layer between OpenCog and the Robokind robot)

The project involves three key components, corresponding to these three researchers:

- The **OpenCog** artificial general intelligence architecture [5,8], an open-source software project initiated by Dr. Goertzel and now developed and maintained by a global open-source community led by Goertzel and others

- The **DeSTIN** framework for machine perception (extensible to action as well), developed by Dr. Itamar Arel and his students at the University of Tennessee, Knoxville; now an open-source software project bundled with OpenCog

- Hanson's robotics hardware, the **Hanson Robokind**, and his open-source software for specifying and maintaining intelligent characters with individual personalities [15],

-- integrated in a manner that is inspired by the prior research and thinking of Goertzel, Hanson and Arel, but also has critical novel aspects. The core of the project is a novel approach to bridging the sensation/action domain as embodied in the Robokind's sensors and actuators and recognized via the (subsymbolic) DeSTIN AI software, and the domain of symbolic, abstract AI cognition as embodied in the OpenCog cognitive architecture.

OpenCog is the most practically advanced software system explicitly oriented toward Artificial General Intelligence. Compared to other contemporary AGI-oriented software systems[2], OpenCog has a more professionally designed and implemented codebase, and a wider range of functionalities. Conceptually, the OpenCog architecture is unique in its integration of a number of powerful learning algorithms within a human-like cognitive architecture.

OpenCog has been used in a variety of commercial projects in areas such as natural language processing, financial analysis and bioinformatics. The main thrust of current OpenCog development is the use of the system to control animated characters in a video-game world – which, following a long tradition in the AI field, is a world where most objects are built from small blocks, whose positions and interactions the AI system can fully understand and manipulate. OpenCog guides characters in the world as they explore and try to achieve their goals. This is a follow-up to earlier work using OpenCog to control virtual pets in a simpler virtual world [17].

OpenCog has also previously been used to control a Nao humanoid robot [18]. However, this was merely a prototyping activity, which was not expected to yield dramatic embodied intelligence, due to OpenCog's lack of a serious perception and action module at that time. The core aim of the present proposal is to remedy this lack.

[2] Samsonovich [16] has given a good general overview of the various AGI software systems at play in the field currently, including e.g. classic systems like SOAR and ACT-R, and more modern ones like (to name just a few) Stan Franklin's LIDA system, Joscha Bach's MicroPsi and Nick Cassimatis's PolyScheme.

***Figure 1**. **Top**: Screenshots of virtual dog controlled by OpenCog engine, as it learns soccer skills via imitation and reinforcement, and builds semantic models of its environment. **Bottom**: Screenshot of "blocks world" being used for current OpenCog experimentation; the robot is one of multiple OpenCog-controlled characters.*

DeSTIN, a machine perception architecture, fills a major gap in OpenCog – the processing of complex, high-bandwidth perceptual data, as is produced by camera eyes or microphones. It possesses a hierarchical architecture similar to the human visual and auditory cortices, and an efficient algorithmic implementation that exploits the massively parallel processing provided by GPU supercomputers. Currently it is being used for recognizing patterns in images and videos, but the architecture can be straightforwardly extended to audition as well, and also beyond perception to handle actuation.

To connect DeSTIN and OpenCog in a maximally effective way, we have designed a unique "semantic perceptual-motor hierarchy", which sits between the two systems, incorporating aspects of each. As depicted in Figure 3, it has DeSTIN's hierarchical structure, but represents knowledge using OpenCog's symbolic semantic network representations, rather than DeSTIN's subsymbolic numerical vectors. This semantic hierarchy naturally maps into both DeSTIN and OpenCog, and enables the two systems to pass information to each other frequently as they operate, enhancing each other's intelligence synergetically.

 

***Figure 2*. *Left*: Dr. Hanson's Robot Einstein, probably the most expressive and emotionally evocative robot face ever constructed. *Right*: Hanson Robokind – the emotionally expressive humanoid robot to be used in the current project, based on a donation of two robots from Hanson Robotics.*

To refine and evaluate this novel approach to bridging OpenCog's symbolic reasoning and DeSTIN's subsymbolic pattern recognition activity, one requires a sufficiently sophisticated platform for receiving sensations and executing actions. The Hanson Robokind robot provides high-quality visual and auditory sensors and servomotor capability, at a level previously available only in research robots costing hundreds of thousands of US dollars).

Furthermore, the Robokind's capability for facial emotional expressiveness is unparalleled in any previous commercially available robot regardless of cost. While not useful for the WPPSI evaluation metric, this aspect is generally valuable in a "childlike AGI" context because it gives the Robokind ability to elicit interesting, informative behavior from humans, thus help the OpenCog engine to gain the emotional and social knowledge it needs to interact effectively with the world.. This facial-expression capability is based on special patented flexible robot skin that was developed in Hanson's prior research robots such as the well-known Robot Einstein, and is brought to the commercial market for the first time in the Robokind (Figure 2).

Project Aims

The high-level objective of the proposed CogBot project is to create a software system instantiating a solution to the most fundamental problem holding back progress toward AGI: the bridging of the symbolic and subsymbolic levels of mental activity. Among the many subproblems to be addressed in this context, the most fundamental regards *intelligent, synergetic interfacing between low-level robotic perception/action data and more abstract AI cognition* – a problem for which we have a novel solution.

Alongside qualitative evaluation and testing of specific technical capabilities, the key measurable aim of the project is to create an OpenCog-powered robot that is able to perform well (above the 90% level) on the WPPSI preschool intelligence test. The specifics of this test are described in the document "The Wechsler Preschool and Primary Scale of Intelligence (WPPSI): A Brief Overview", provided as a companion to this proposal; and some issues regarding the application of this test in a robotics context are described in the document "Utilizing Child IQ Tests to Measure Robot Intelligence", also provided. As noted there, the critical intelligent capabilities required for success on the WPPSI test are:

1. **Natural language question answering** (regarding information generally known to young children, and the immediate physical situation of the robot)
2. **Object, event and part identification** (especially for objects depicted in pictures, and commonly known to young children)
3. **Object manipulation** (minimally of objects sitting on, and slid around on, a tabletop)
4. **Visual pattern recognition** (of patterns in drawing, and textures on and shapes of physical objects)

5. **Simple drawing** (not necessarily of pictures nor words, but lines and other simple symbols, using pencil or marker on paper)
6. **Instruction following** (regarding tasks involved in the intelligence test, but also more broadly)
7. **Pragmatic interaction regarding task assignment** (because the robot will not be fed the test questions in an artificial way, but it will rather be posed them as part of a general, everyday robot-interaction)

On a technical level, the key aims of the project in pursuit of these capabilities are:

1. To create an intermediate "semantic hierarchy" connecting DeSTIN and OpenCog
2. To make various changes identified as necessary to DeSTIN, to enable effective connection with the semantic hierarchy
3. To adjust OpenCog's internal cognitive algorithms for optimal functionality on the data coming from the semantic hierarchy (mainly the PLN probabilistic logic engine, the conceptual blending algorithm for creating new concepts, and the Fishgram pattern recognition algorithm)
4. To integrate the combined OpenCog/DeSTIN software with Hanson's robot control software, and with open-source text-to speech (Festival) and speech-to-text (Sphinx) engines, to enable the system to control a Robokind humanoid robot
5. To refine and evaluate this robot's capability to carry out simple "child-like" tasks in a robot lab furnished with appropriate preschool-like equipment

Items 4 and 5 will focus mainly but not exclusively on tabletop-interaction based tasks, as those are the ones that WPPSI focuses on. However, we will not create a robot that can only interact or display intelligence across a tabletop!

Strategy and Methodology

The fulfillment of the project's aims involvse a significant amount of highly technical software development, which will be carried out by the two proposed new hires, together with the 6 junior members of the current OpenCog Hong Kong team. Figure 3 gives a high-level diagram illustrating the integration of DeSTIN and OpenCog, in the context of humanoid robotics. Figure 4 depicts more of the internals of the OpenCog architecture, as currently being used for animated agent control. From OpenCog's perspective, the robot and game character are essentially the same sort of entity – the difference lying in the critical "symbolic/subsymbolic converter" component that translates between the language of perception/action (DeSTIN states) and the language of abstract cognition (OpenCog's internal Atom representation). Figure 5 gives a deeper view of OpenCog/DeSTIN integration.

Note that in this design, the bulk of the robot's "mind" lives on a laptop or PC assumed to be on the same wifi network as the robot. In a commercial product based on this design, a portion of the robot's mind may also live in the cloud, allowing sharing of knowledge between different robots. The processor on board the robot handles real-time action response, low-level sensory processing, and communication with the rest of the robot's mind via wifi.

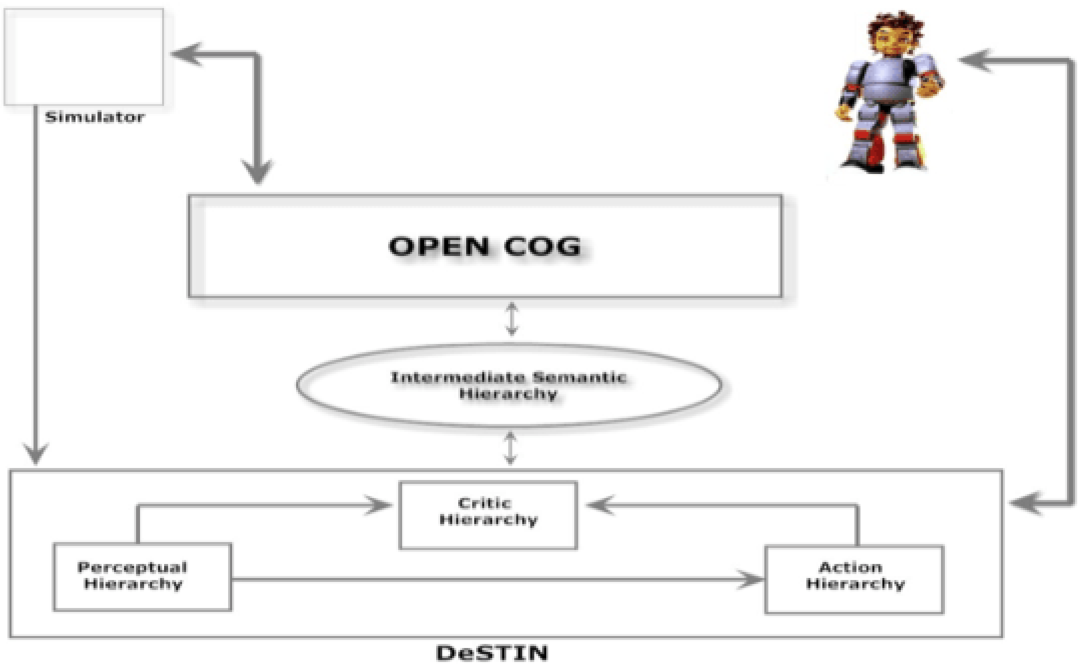

Figure 3. *High-Level Architecture of proposed system. Key components include the Hanson Robokind robot, the OpenCog cognition engine, the DeSTIN perception/action engine, and a novel "symbolic/subsymbolic converter" translating between DeSTIN's subsymbolic perception/action language and OpenCog's semantic-network language.*

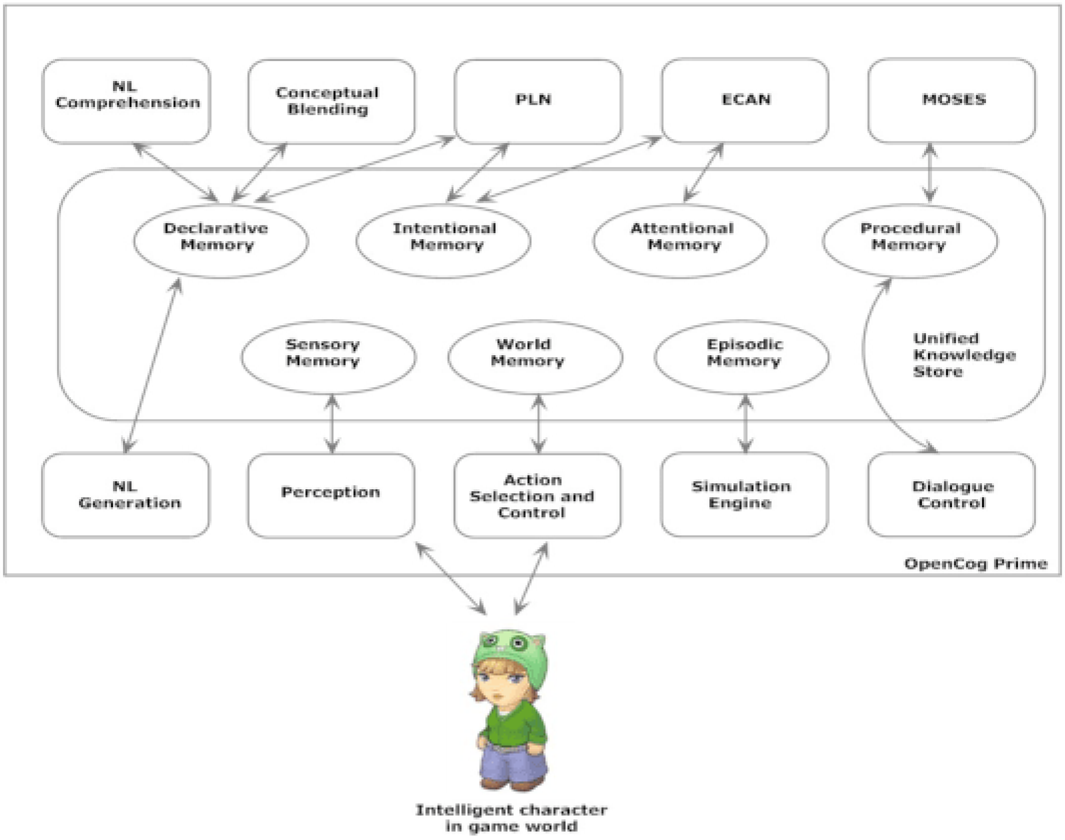

Figure 4*: Key components of the OpenCog integrated artificial general intelligence architecture, shown in the context of intelligent game character control. The proposed robotic application is similar, but with the DeSTIN perception/action component used as an intermediary between OpenCog and the robot, as robot perception/action is much subtler than its analogue in the virtual world.*

EFTA01103515

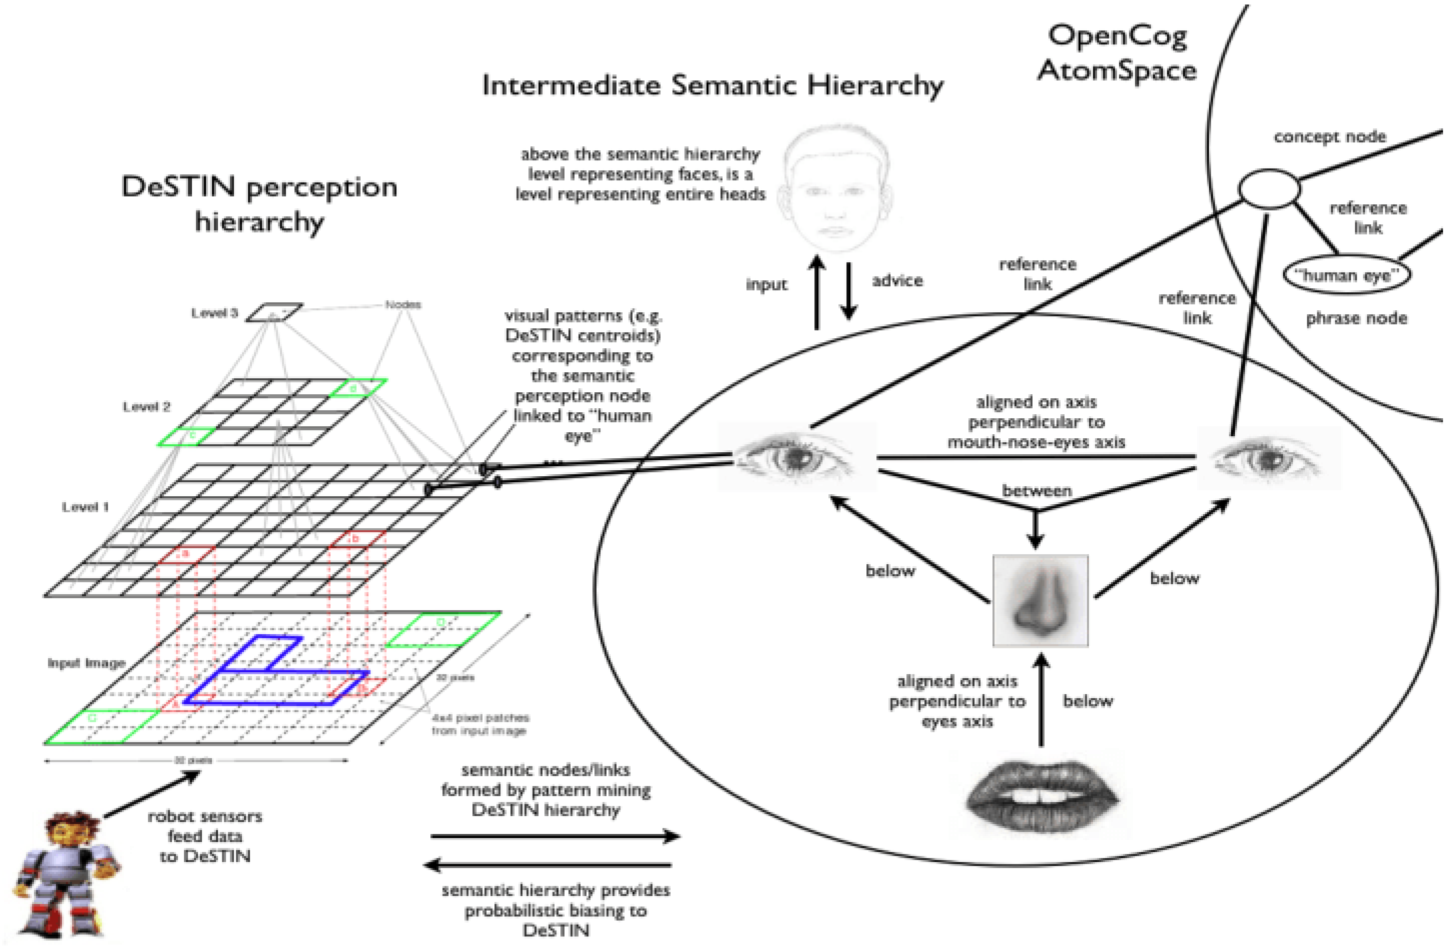

***Figure 5**: An in-depth illustration of the "intermediate semantic hierarchy" referenced in Figure 3. In the context of face recognition, this shows the interfacing between DeSTIN and OpenCog's "cognitive semantic network" knowledge representation, via means of a symbolic-subsymbolic translation layer that utilizes OpenCog's knowledge representation atoms but possesses DeSTIN's hierarchical structure. This unique interfacing layer is the central scientific novelty of the proposed project.*

Anticipated Outcomes and Impacts

This project offers the potential to introduce a dramatic advance in artificial general intelligence and cognitive robotics. It will serve as a platform for further dramatic advances, in the form of ongoing [redacted] aimed at producing software enabling humanoid robots to achieve humanlike general intelligence beyond the early childhood level. Its successful completion is expected to have broad impact on the AI field, inspiring other researchers to pursue integrated cognitive architectures for intelligent agent control, and in general helping to revive research interest in the original, ambitious goals of the AI field.

Concretely, in order to ensure this influence happens, we intend to publish at least three papers in high-impact journals summarizing the results of the project, along with making a highly-publicized release of the open-source code developed, and launching a series of YouTube videos showing the intelligent robot children in action. Presentations at appropriate conferences will also be done, including academic robotics and AI conferences (AAAI, IEEE) as well as futurist conferences such as the Singularity Summit, Humanity+, World Future Society, etc.

While the current proposal is focused on pure research, we also intend to encourage the commercialization of the technology developed in the project, initially via collaboration with Hanson Robotics and Jetta (the Hong Kong firm that manufactures Hanson's robots) on the creation and marketing of intelligent robot toys. Hanson Robotics has carefully explored the business and technical aspects of this sort of product offering, but currently lacks the AI technology to make it work on their own.

Of the $21 billion market for remote-control electronic toys, about $200 million consists of higher-end robotic devices costing $100 or more (including e.g. the Nao, smart drones like the Parrot, and the Kondo robot kits).[3] At lower price points, it's hard to draw the line between robot toys like RoboSapien and "mere" electronic remote control devices. Clearly the market for ~$100 robot toys could be grown dramatically via the entrance of radically superior technology into the marketplace. One could imagine a commercial toy robot with an initial hardware cost of ~$99, plus a monthly subscription fee of perhaps $4.99, the latter buying online access to OpenCog AGI servers supplying the robots with advanced intelligence.

Beyond this, there are a number of possibilities for commercialization of the specific cognitive robotics technology expected to result from this project, including in:

- **The robotics industry.** The advances in general intelligence developed for robots will make them suitable to take on more roles in the service industry, among others.
- **The toy industry.** With toys becoming increasingly sophisticated and electronic, our proposed software can be applied to drive the behavior of new electronic toys.
- **The consumer electronics industry.** Consumer electronics such as tablet computers and smart phones could benefit greatly from being more intelligent, more responsive to users and more predictive of users' needs.

And of course, the follow-on AGI development enabled by success in our project could have much broader commercial impact.

We are also enthused about the broader social implications of advanced artificial general intelligence making its initial advent in the guise of friendly, childlike humanoid robots. In the future as AGI advances beyond the childlike level, human attitudes toward AGI may become complex and contentious; we believe the best path to an agreeable future is one in which people and early-stage AGIs have a relationship of mutual understanding.

Budget

Please see separate budget spreadsheet.

Timeline/Milestones

Please see separate document

References

1. B. Goertzel and C. Pennachin, Artificial General Intelligence. Springer, 2005.
2. [AGI-01] B. G. Stan Franklin and P. W. (Editors), Proceedings of First Conference on Artificial General Intelligence. IOS Press, 2008.
3. [AGI-02] P. H. Marcus Hutter and B. G. (Editors), Proceedings of Second Conference on Artificial General Intelligence. Atlantis Press, 2008.
4. N. C. (Editor), Human-Level Intelligence (Special Issue of Artificial Intelligence Magazine). AAAI, 2006.
5. B. Goertzel, "Opencog prime: A cognitive synergy based architecture for embodied artificial general intelligence," in Proceedings of ICCI-09, Hong Kong, 2009.
6. Arel, D. Rose, and T. Karnowski, "A deep learning architecture comprising homogeneous cortical circuits for scalable spatiotemporal pattern inference." NIPS 2009 Workshop on Deep Learning for Speech Recognition and Related Applications, Dec 2009.
7. Thomas P. Karnowski, Itamar Arel, Derek Rose: Deep Spatiotemporal Feature Learning with Application to Image Classification. ICMLA 2010: 883-888
8. Goertzel, Ben (2011). Integrating a Compositional Spatiotemporal Deep Learning Network with Symbolic Representation/Reasoning within an Integrative Cognitive Architecture via an Intermediary Semantic Network. Proceedings of AAAI Symposium on Cognitive Systems, Arlington

[3] Figures via personal communication from Mark Tilden, founder & chief scientist of Wowee, maker of RoboSapien

VA
9. Goertzel, Ben and Hugo de Garis. XIA-MAN: An Integrative, Extensible Architecture for Intelligent Humanoid Robotics. AAAI Symposium on Biologically-Inspired Cognitive Architectures, Washington DC, November 2008
10. Goertzel, Ben and Stephan Vladimir Bugaj. AGI Preschool: A Framework for Evaluating Early-Stage Human-like AGIs. *Proceedings of the Second Conference on Artificial General Intelligence*, Atlantis Press.
11. Sam Adams, Itmar Arel, Joscha Bach, Robert Coop, Rod Furlan, Ben Goertzel, J. Storrs Hall, Alexei Samsonovich, Matthias Scheutz, Matthew Schlesinger, Stuart C. Shapiro, John Sowa (2012). Mapping the Landscape of Human-Level Artificial General Intelligence. AI Magazine, Winter 2012.
12. P. Hayes and K. Ford., "Turing test considered harmful," IJCAI-14, 1995
13. Kurzweil, Ray (2005). The Singularity Is Near. Viking.
14. Breazeal, Cynthia (2002). Designing Sociable Robots. MIT Press.
15. Hanson D. "Bioinspired Robotics", chapter 16 in the book Biomimetics, ed. Yoseph Bar-Cohen, CRC Press, October 2005.
16. Hanson, David and V. White. (2004). "Converging the Capabilities of ElectroActive Polymer Artificial Muscles and the Requirements of Bio-inspired Robotics", Proc. SPIE's Electroactive Polymer Actuators and Devices Conf., San Diego
17. Samsonovich, A. V. (2010). Toward a unified catalog of implemented cognitive architectures (review). In Samsonovich, A. V., Jóhannsdóttir, K. R., Chella, A., and Goertzel, B. (Eds.). *Biologically Inspired Cognitive Architectures 2010: Proceedings of the First Annual Meeting of the BICA Society. Frontiers in Artificial Intelligence and Applications,* vol. 221, pp. 195-244. Amsterdam, The Netherlands: IOS Press. ISSN 0922-6389.
18. Goertzel, Ben, Cassio Pennachin, Nil Geissweiller, Moshe Looks, Andre Senna, Ari Heljakka, Welter Silva, Carlos Lopes . An Integrative Methodology for Teaching Embodied Non-Linguistic Agents, Applied to Virtual Animals in Second Life, in *Proceedings of the First AGI Conference*, Ed. Wang et al, IOS Press