

# OpenCog Hong Kong Project Interim Report, May 2011

*Prepared for the Jeffrey Epstein Foundation*

## Contents

- Background
- Articulation of Specific Goals
- Technical Achievements So Far
- Staffing
- Screenshots of In-Development Game World

## Background

The OpenCog Hong Kong project is a collaboration between Hong Kong Polytechnic University and Novamente LLC, aimed at the use of the OpenCog AI engine to create a working prototype of simple “artificial general intelligence” in a video game world. The prototype is supposed to take the form of an open-source software toolkit that can be utilized and extended, and the principal aim of the project is to provide a concrete context in which to implement and test various aspects of the OpenCog system.

The project is funded 90% by the Hong Kong government and 10% by Novamente LLC. Novamente LLC’s portion amounts to roughly US\$40K altogether. The initial \$20K of Novamente’s portion was funded by a donation from Humanity+, which in turn was funded by a donation to Humanity+ from the Jeffrey Epstein Foundation.

The funding is earmarked to pay primarily for staff salaries: one AI PhD to lead the team, plus 3 other AI programmers, one game programmer, and one artist.

The project formally began at the start of October 2010, but in fact staff were brought on more slowly, due to visa issues (most staff were brought to Hong Kong from overseas, due to the desire to work with staff having prior OpenCog experience). A detailed rundown of project staffing is given later in this document. A quick summary is that the current staff have been on board from between 4 and 6 months, at this point (May 2011).

It was originally thought the second \$20K of Novamente’s contribution would be due in September 2011, but recently it’s been revealed that it’s actually due before the end of May 2011. This is due to Hong Kong government regulations, and is not negotiable.

## Articulation of Specific Goals

One of the things achieved during the first months of the project has been to more clearly articulate the project goals. (A lot of technical work was also achieved; see the following section for details.)

We have chosen a specific game world to use for testing and demo-ing, though the software built will apply in essentially any game world. The game world chosen is inspired by the game Minecraft, but is implemented in the Unity3D game engine (which provides a very flexible API for integration of OpenCog or other AI systems). Some images of Minecraft and our in-progress game world are given at the end of this document. A decision was also made to have the AI control a demo character taking the appearance of a robot (although the robot artwork is not ready yet, so that the screenshots at the end of the document show instead the "woman in a red dress" imported from earlier OpenCog demos).

The basic goals of the OpenCog Hong Kong project, in this context, are for OpenCog to control an artificial agent that can:

- Build structures with blocks so as to achieve its goals in the game world
- Use simple English to describe what it's doing and explain why, and ask questions relevant to its goals

Simple examples would be:

- Learning to build steps or ladders to get desired objects that are high up
- Learning to build a shelter to protect itself from aggressors
- Learning to build structures resembling structures that it's shown (even if the available materials are a bit different)
- Learning how to build bridges to cross chasms

Of course, the AI significance of learning tasks like this all depends on what kind of feedback the system is given, and how complex its environment is. It would be possible to do things like this in a trivial and highly specialized way, but that is not the intent of the project – the goal is to have the system learn to carry out tasks like this using general learning mechanisms and a general cognitive architecture, based on embodied experience and only scant feedback from human teachers. This will provide an outstanding platform for ongoing AGI development, as well as a visually appealing and immediately meaningful demo for OpenCog.

We have identified some specific tasks that ██████ like to see the system execute by August 2011, though we're unsure if this will be realistic (it might take a month or two more). An earlier version of these tasks are depicted in the document Bananas\_demo.doc. The vision of the demo has changed a bit since then; for

instance, in the new version, instead of seeking bananas, the character will look for batteries. But the essence remains the same. The core ideas are that the AI character should be able to, e.g.

- Watch another character build steps to reach a high-up object
- Figure out via imitation of this that, in a different context, building steps to reach a high up object may be a good idea
- Also figure out that, if it wants a certain high-up object but there are no materials for building steps available, finding some other way to get elevated will be a good idea that may help it get the object

With a bit of luck we will be able to demo this at the AGI-11 conference on Google campus in early August 2011.

### **Technical Achievements So Far**

The OpenCog Hong Kong project has led to a large amount of work on OpenCog getting done, both system engineering and AI advances. There have been too many small tasks completed for it to be sensible to list them all here, but here are some of the major achievements.

A roadmap for the whole OpenCog HK project is attached as a separate document. Here I just give an informal run-down of what we've done so far. A big-picture roadmap for the whole OpenCog project is also attached, including the HK project as a subset.

On the AI side:

- Implementation of the **Psi model of emotion and motivation** (created by German psychologist Dietrich Dorner and refined by Joscha Bach in his MicroPsi AGI system) inside OpenCog, yielding the OpenPsi system. This is a complete overhaul of OpenCog's motivation and goal system. (Note that Joscha Bach visited Hong Kong from Germany in April 2011 to help with this work.)
- Implementation of "**frequent subgraph mining**" in OpenCog, for identifying frequently occurring patterns in the game world and in the system's mind – for instance "when I plug myself in I no longer need electricity", "chairs and tables are often found near each other", "Bob often hits me", etc. This was done via integrating the **SUBDUE graph mining software** with OpenCog
- Implementation of a **probabilistic planner** based on OpenCog's Probabilistic Logic Networks reasoning system, enabling simple planning of actions in the game world
- Implementation of a fairly **complete virtual-world perception system** that registers events observed in the game world as logical relationships in OpenCog's knowledge base. Such a system was present in OpenCog

previously, but only in a very simplified form, not adequate for the current project.

And on the system engineering side:

- Creation of a framework enabling **OpenCog MindAgents** (cognitive processes) to be coded **in Python as well as C++**. As Python is an easier and faster language to program in, this should speed up OpenCog software development, as well as make it accessible to a larger pool of programmers.
- **Integration of OpenCog with Unity3D**, a more powerful and flexible game engine than the Multiverse engine that OpenCog was integrated with before
- **Customization of the MineCraft-like plug-in** for Unity3D, to make it suitable as a world for OpenCog demos

Note also that

- Three OpenCog related papers were accepted for the AGI-11 conference: one on OpenPsi, one on information geometry for attention allocation, and one on OpenCog/DeSTIN integration for computer vision (the latter two not directly based on the Hong Kong project's work)
- Two OpenCog related papers were accepted for the AAAI 2011 conference: one on MOSES-PLN integration, and another on lifelong learning

### Staffing

- **Dr. Joel Pitt**, team leader (from New Zealand), started November 2010
  - As well as managing the team, Joel has been doing various AI and game programming tasks for the project
- **Jared Wigmore** (also from New Zealand), started toward the end of December 2010
  - Working on probabilistic inference and pattern mining in OpenCog
- **Troy Huang and Zhenhua Cai** (from Xiamen University in mainland China) started in January 2011. Note: *They are exchange graduate students and will return to Xiamen in January 2012. Zhenhua will return to Hong Kong in September 2012 if there is ongoing funding available.*
  - Zhenhua is working on OpenPsi, an implementation of the Psi model of emotion and motivation inside OpenCog
  - Troy is working on the practicalities of connecting OpenCog to the game world, and on emotion modeling inside OpenCog
- **Cord Krohn** started in November 2010, working on the graphic arts aspect
- **Lester Lam** started in November 2010, working on game programming, and left the project in April 2011

Cord and Lester's salaries were covered by additional Hong Kong government funding, separate from the funding mentioned above.

Two additional team members have been made job offers and the paperwork to bring them on is pending:

- **Shujing Ke**, an AI PhD student from Xiamen, with 3 years commercial game programming experience
  - Shujing will take charge of applying imitation and reinforcement learning aspects of OpenCog to the appropriate context, and also help with game programming
- **Dr. Ruiting Lian** (PhD to be completed in fall 2011), a computational linguistics expert from Xiamen
  - Ruiting will re-factor the OpenCog language comprehension and generation system to make it more adaptable based on an agent's experience in a world, i.e. to enable simple language learning
  - She is the original author of OpenCog's language generation system

Potential additional team members are being interviewed.

## Screenshots

Minecraft itself, the game that inspired our test/demo game world for the project:



On the next page, some shots of our current, crude Minecraft-like world in Unity3D. Note that we will soon replace the woman with a robot!

