

Exploring Android Developmental Psychology in a Simulation World

Ben Goertzel (ben@goertzel.org)

Applied Research Lab for National and Homeland Security, Virginia Tech, 2000 N. 15t St., Ste. 50
Arlington, VA 22216 USA

Ari Heljakka (heljakka@iki.fi), Stephan Vladimir Bugaj (stephan@bugaj.com),

Cassio Pennachin (cassio@vettalabs.com)

Novamente LLC, 1405 Bernerd Place
Rockville, MD 20851 USA

Moshe Looks (moshe@metacog.org)

Department of Computer Science, Washington University, One Brookings Drive
St. Louis, MO 63130 USA

Abstract

A research program oriented toward exploring the development of android cognition in the context of a 3D simulation world is described. The simulation world itself, AGISim, is an open-source project built on a game engine, which allows human-controlled and AI-controlled android agents to interact inside a simulated apartment. AGISim has been used for the embodiment of android agents controlled by the Novamente AI Engine, in the context of an AI-education program guided by Piagetan learning theory. Recent experiments have focused on teaching these android agents to understand the notion of the permanent identity of an agent or an object. Experiments involving the spontaneous learning of “theory of mind” have been designed and will be performed in the near future.

Introduction

The main stream of research in android science focuses, as it should, on the engineering, programming and instruction of physical android robots. However, the current state of android hardware is still relatively primitive, which means that in practical research, cognition tends to get short shrift, since achieving effective android perception and action are still such major obstacles. Thus, we suggest that it is worthwhile to complement work on physical android robotics with work in computer simulation worlds, in which human-controlled simulated androids interact with AI-controlled androids in simulated environments.

No simulation world running on currently affordable hardware will be able to provide a fully accurate simulation of the perceptual and motor-control challenges faced by physical android robots. However, we suggest that contemporary simulation worlds, appropriately utilized, can nonetheless permit effective simulation of many of the cognitive challenges that physical android robots face.

With this philosophy in mind, we have created a 3D simulation world called AGISim (Goertzel et al, 2006), and begun using it to teach an AI system to control a simulated android, in the context of interactions with a human-controlled simulated android. Within this framework we are pursuing an AI-teaching program loosely guided by Piagetan developmental psychology. Our current focus is

on infant-level cognition such as the understanding of the permanence of objects and agents. The next phase of teaching will focus on “theory of mind” – on encouraging the AI system to come to its own understanding of the intentions and beliefs and knowledge of other cognitive agents, based on its interactions with them in the simulated world.

AGISim and Novamente

The AGISim simulated world is being developed as an open-source project¹, led by the first two authors, and is based on the CrystalSpace² 3D game engine, which may be configured to display realistic physics. It allows AI systems and humans to control android agents, and to experience the simulated world via multiple senses, as well as having the capability to chat with each other directly through text.

It is intended that the experience of an AGI controlling an agent in AGISim should display the main qualitative properties of a human controlling their body in the physical world. The simulated world should support the integration of perception, action and cognition in a unified learning loop. And, it should support the integration of information from a number of different senses, all reporting different aspects of a common world. With these goals in mind, we have created the initial version of AGISim as a basic 3D simulation of the interior of a building, with simulations of sight, sound, smell and taste. An agent in AGISim has a certain amount of energy, and can move around and pick up objects and build things. While not an exact simulation of any specific physical robot, the android agent an AI controls in AGISim is designed to bear sufficient resemblance to a simple humanoid robot that the porting of control routines learned in AGISim to a physical robot should be feasible, though certainly not trivial.

Our work with AGISim to date has focused on controlling android agents in AGISim using the Novamente AI Engine (or NAIE; Goertzel, 2006; Looks, Goertzel and Pennachin, 2004), a comprehensive AI architecture that synthesizes perception, action, abstract cognition, linguistic

¹ sourceforge.net/projects/agisim

² crystal.sourceforge.net

capability, short and long term memory and other aspects of intelligence, in a manner inspired by complex systems science. Its design is based on a common mathematical foundation spanning all these aspects, which draws on probability theory and algorithmic information theory among other areas. Unlike most contemporary AI projects, it is specifically oriented towards artificial *general* intelligence (AGI), rather than being restricted by design to one narrow domain or range of cognitive functions.

The NAIE integrates aspects of prior AI projects and approaches, including symbolic, neural-network, evolutionary programming and reinforcement learning. However, its overall architecture is unique, drawing on system-theoretic ideas regarding complex mental dynamics and associated emergent patterns. The existing codebase implements roughly 60% of the overall design, and is being applied in bioinformatics, NLP and other domains.

Cognitive Development in Simulated Androids

Jean Piaget, in his classic studies of developmental psychology (Inhelder and Piaget, 1958), conceived of child development as falling into four stages, each roughly identified with an age group: infantile, preoperational, concrete operational, and formal. While Piaget's approach is out-of-date in some ways, recent researchers have still found it useful for structuring work in computational developmental psychology (Shultz, 2003); we have modified the Piagetan approach somewhat for usage in our own work. The basic Piagetan stages are as follows:

- *Infantile*: Imitation, repetition, association. Object permanence – infants learn that objects persist even when not being observed.
- *Preoperational*: Abstract mental representations. Word-object and image-object associations become systematic rather than occasional. Simple syntax.
- *Concrete*: Abstract logical thought applied to the physical world: conservation laws; more sophisticated classification; theory of mind – an understanding of the distinction between what I know and what others know. Classification becomes subtler.
- *Formal*: Abstract deductive reasoning contextually and pragmatically applied, the process of forming then testing hypotheses, etc.

We have carried out learning experiments involving the NAIE and AGISim, corresponding to aspects of Piaget's early stages of development.

We have shown that a Novamente-powered simulated android can learn, via its interactions with human-controlled simulated android agents, to carry out basic cognitive tasks like word-object association and understanding the permanence of objects and agents. In particular, for object permanence, we have created a simulation of the "A-not-B" task commonly used in developmental psychology (see e.g. Thelen and Smith, 1994), and shown that Novamente's ability to solve this task is specifically tied to its correct use of a specific inference rule called the "Rule of Choice."

The teacher hides an object in location A repeatedly, then eventually hides it in location B and asks the AI agent to find it. Human babies less than 9 months of age will often look in location A, but older babies look in B. The NAIE learns through interactive experience to look in location B – it learns that objects exist even when unobserved.

We have also shown that a Novamente-powered simulated android can successfully classify objects in the simulation world, at a level corresponding roughly to Piaget's preoperational phase. A typical experiment involves distinguishing chairs from couches from boxes based on their appearance in various contexts.

The next stage of our work will involve Piaget's concrete operational phase, and in particular "theory of mind." We aim to teach a Novamente-controlled simulated android agent to understand that other simulated android agents possess knowledge and beliefs separate from their own. It is important to stress that in our approach this knowledge is not programmed: it is learned based on embodied social interaction. To teach "theory of mind," for example, the teacher can hide an object in such a way that Novamente can see that it sees the hiding action but its playmate cannot; and then it can be checked if Novamente can predict where the playmate will look for the object. Without adequate theory of mind, Novamente will predict the playmate will look in the place the object was hidden. With adequate theory of mind, Novamente will predict the playmate will search for the object in the most obvious place.

The knowledge obtained via this sort of experiment is not particular to the AGISim simulation world, but is of generic value, and should be largely portable to the domain of physical android robotics, at such time when the NAIE is used to control a physical android robot.

References

- Goertzel, B. (2006). Patterns, Hypergraphs and Artificial General Intelligence. *Proceedings of IJCNN 2006, Vancouver CA*
- Goertzel, B., M. Looks, A. Heljakka, and C. Pennachin (2006). Toward a Pragmatic Understanding of the Cognitive Underpinnings of Symbol Grounding, in *Semiotics and Intelligent Systems Development*, Edited by Ricardo Gudwin and João Queiroz
- Inhelder, B. and J. Piaget. (1958). *The Growth of Logical Thinking from Childhood to Adolescence*. New York: Basic Books.
- Looks, M., B. Goertzel and C. Pennachin. (2004). "Novamente: an integrative architecture for Artificial General Intelligence." *Proceedings of AAAI 2004 Symposium on Achieving Human-Level AI via Integrated Systems and Research*, Washington DC
- Shultz, T. (2003). *Computational Developmental Psychology*. Cambridge, MA: MIT Press
- Thelen, E. and L. Smith (1994). *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA: MIT Press