

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 programme 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 Piagetian 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 Piagetian developmental psychology. Our current focus is

on infant-level cognition such as the understanding of the permanence of objects and agents, exemplified for instance by Piaget’s A-not-B task. 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 probabilistic inference, 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 code base 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). As will be described below in some detail, we have modified the Piagetian approach somewhat for usage in our own work.

The basic Piagetian 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. A little later we will discuss a particular example related to 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 the NAIE's ability to solve this task

is specifically tied to its correct use of a specific inference rule called the "Rule of Choice."

We have also shown that an NAIE-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.

Piagetian Stages and Uncertain Inference

To address the various critiques of Piaget's theories that have been made (Commons et al, 1982, 1988; Pascual-Leone and Smith, 1989; Fischer, 1980; Marchand, 2001), and to better bridge the gap between the human developmental psychology on which Piaget's theory was based and the non-human cognitive structures within the NAIE, we have created a novel theory of cognitive developmental stages, defined in terms of the control of uncertain inference trajectories. Our theory of developmental stages is oriented specifically toward AI systems like Novamente that are founded on uncertain logical inference rules "control schemata" (sometimes learned, sometimes pre-programmed) that determine the order in which inference rules are to be applied. It may be that our modified version of Piagetian theory also has applicability to human psychology, but at the moment we are focusing on its applications to AI and humanoid robotics.

The stages in our theory are defined as follows:

1. *Infantile*: Able to recognize patterns in and conduct inferences about the world, but only using simplistic hard-wired (not experientially learned) inference control schemata, along with pre-heuristic pattern mining of experiential data.
2. *Concrete*: Able to carry out more complex chains of reasoning regarding the world, via using inference control schemata that adapt behavior based on experience (reasoning about a given case in a manner similar to prior cases).
3. *Formal*: Able to carry out arbitrarily complex inferences (constrained only by computational resources) via including inference control as an explicit subject of abstract learning.
4. *Reflexive*: Capable of self-modification of internal structures. (In the case of the NAIE, this process is very direct and thorough: one key architectural difference between humans and AI's is the latter's vastly greater capability for self-modification)

In this approach Piaget's preoperational phase appears as transitional between the infantile and concrete operational phases; and, following a number of recent thinkers, we have explicitly introduced a post-formal stage as well.

The semantics of our stages is similar but not identical to Piaget's. Our stages are defined via internal cognitive mechanisms, which represent not only abstract knowledge, but also perceptual information, and both cognitive and

operational task skills. Uncertain inference is assumed to allow the developing cognitive system to reason through experience in a fuzzy and context-variant way, rather than requiring a fully-formed absolute logical formulation of each particular situation.

We posit that these developmental stages correspond to the ability to solve certain classes of problems in a generalizable way. For instance, we suggest that it is only through inference control schemata which adapt based on experience that uncertain inference-based AI systems can learn to consistently solve Piagetian concrete-operational tasks in a way that provides knowledge suitable for further generalization. Of course, it may be that minds using hard-wired inference control schemata (typical of the infantile stage) can still solve some Piagetian concrete-operational tasks. Such brittle approaches to solving such tasks, historically, have proved unable to generalize sufficiently and resulted in permanently “brittle” AI systems which are very limited in capability.

We have designed an AGISim based learning programmed for the NAIE based on these stages, while at the same time accounting for key ideas of dynamical developmental psychology:

1. Not all tasks will develop into the same stage at the same time.
2. Stages represent not only abstract cognitive abilities, but also interrelated perceptual and operational abilities.

Since development is a dynamical feedback system between perception, action, and abstraction, different abstract cognitive tasks may develop into different stages at different times based on the amount of experience with relevant actions and perceptions.

The A-not-B Error

In Piaget's classic "A-not-B error," the teacher hides an object in location A repeatedly, then eventually hides it in location B and asks the subject (a baby or in our case, an AI agent) to find it (Thelen and Smith, 1994; see also Spencer et al, 2001).. Human babies less than 9 months of age who have successfully uncovered a toy at location A in prior trials will often continue to reach to that location even after they watch the toy hidden in a nearby location B. Older babies will look in B after they've seen the toy hidden in B. In some cases, infants will look at B (the correct location) but reach for A (the incorrect location), indicating a complex and only semi-consistent internal knowledge base.

We have created a simulated-humanoid-robotics emulation of this Piagetian scenario in AGISim, and used the NAIE in this context. The simulation involves a humanoid agent controlled by the NAIE, and a humanoid teacher agent controlled by a human. The A-not-B task is represented via the placement of a toy bunny in boxes which the agents may open or close by pushing buttons. Thus, rather than an abstract mathematical problem, the NAIE-controlled agent is presented with a problem involving

integrated perception, action and cognition, analogous to the problems presented by human infants.

What we find is that, like older babies, the NAIE learns through interactive experience to look in location B – it learns that objects exist even when unobserved. However, it is also possible to emulate younger-baby behavior within the NAIE by modifying the way its inference rules operate.

The NAIE system's ability to solve the A-not-B task correctly is specifically tied to its correct use of a specific inference rule called the “Rule of Choice.” This is the rule that allows the system's inference engine to correctly choose between two competing estimates of the truth value of the same relationship. In this case, one estimate comes from the simple repetitive fact that the toy has often been found in location A before; and another estimate comes from the knowledge that inanimate objects tend to be found where they were last seen. The latter knowledge is more general and a properly functioning Rule of Choice will choose it over the former in this case.

We can also emulate, in the NAIE, the infantile behavior in which the A-not-B error is made via reaching but not via looking. This occurs when implications joining sensations to actions are made and revised directly without the intervention of more abstract representations that abstract away from particular sensation and action modalities. This may happen in the NAIE when the parameters of the concept-formation and inference components are not tuned to encourage the learning of abstractions.

The inferences involved in the A-not-B task are “infantile” in the sense of our above developmental theory, in the sense that they can be carried out using simple non-adaptive forward-chaining or backward-chaining inference control. The difference between incorrect and correct behavior on this task has to do, not with the high-level properties of inference control schemata, but with the correct usage of a particular inference rule. However, it is the execution of specific inferences like this one from which general inference control patterns are learned, enabling adaptive inference control as is required by later developmental stages.

Conclusion

We have described a research programmed aimed at teaching an AI system to control a humanoid agent in a simulated environment, and discussed some of the early steps we have taken along this path.

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. And the lessons learned in doing this work should be in large measure extensible beyond the NAIE to AI-driven humanoid robotics in general.

References

- Commons, M., F. Richards & D. Kuhn. (1982). Systematic and metacognitive reasoning: a case for a level of reasoning beyond Piaget’s formal operations. *Child Development*, 53, 1058-1069.
- Commons, M., Trudeau, E.J., Stein, S.A., Richards, F.A., & Krause, S.R. (1998). Hierarchical complexity of tasks shows the existence of developmental stages. *Developmental Review*, 18, 237-278.
- Fischer, K. (1980). A theory of cognitive development: control and construction of hierarchies of skills. *Psychological Review*. 87, 477-531.
- Goertzel, B. (2006). Patterns, hypergraphs and artificial general intelligence. *Proceedings of IJCNN 2006, Vancouver CA*.
- Goertzel, B., M. Looks, A. Heljakka, & 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, & 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.
- Marchand, H. (2001). Reflections on PostFormal thought. In *The Genetic Epistemologist*, 29(3).
- Pascual-Leone, J. and J. Smith. (1969). The encoding and decoding of symbols by children: a new experimental paradigm and neo-Piagetian model. *Journal of Experimental Child Psychology*, 8, 328-355.
- Shultz, T. (2003). *Computational developmental psychology*. Cambridge, MA: MIT Press.
- Thelen, E. & L. Smith (1994). A dynamic systems approach to the development of cognition and action. Cambridge, MA: MIT Press.
- Spencer, J.P., L.B. Smith, and E. Thelen. (2001). Biobehavioral development, perception, and action tests of a dynamic systems account of the A-not-B error: The influence of prior experience on the spatial memory abilities of two-year-olds. *Child Development*, vol. 72 issue 5, Sept/Oct 2001, 1327.
- Thelen, E., G. Schoner, C. Scheier, and L.B. Smith. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral Brain Science*, 24(1), 1-34.