Developmental Robotics

AAAI Spring Symposium
Stanford University
Palo Alto, California
21-23 March 2005

Pierre.Dangauthier@imag.fr

Abstract:
This document presents some personal notes on the Developmental Robotics symposium held in Stanford and organized by the American Association for Artificial Intelligence. It's not intended to be exhaustive nor well organized, but it's rather a mix between my personal thoughts, key ideas, discussed issues and quotations related to Developmental Robotics. I have a hard copy of the presented papers.

General information

This Workshop was composed of 8 symposia:

- AI Technologies for Homeland Security
- Challenges to Decision Support in a Changing World
- Developmental Robotics
- Dialogical Robots: Verbal Interaction with Embodied Agents
- Knowledge Collection from Volunteer Contributors (KCVC05)
- Metacognition in Computation
- Persistent Assistants: Living and Working with AI
- Reasoning with Mental and External Diagrams.

The Developmental Robotics Symposium was organized by Douglas Blank and Lisa Meeden and was constituted by several papers sessions, posters and panel sessions.
Web Site: http://www.cs.brynmawr.edu/DevRob05/
Related conferences:

- International Conference on Development and Learning
- Epigenetic Robotics 2

Community: A common wish is to set up a community around DevRob, by several ways:

- A WIKI should be put on-line soon
- Doug Blank recommends to use common simulation tools
  - PYRO (python robot programming, easy to use for teaching and research) linked with:
  - Player/Stage/gazebo: Open source Robotic simulator platform, with dynamics, 2D, 3D…
- A website presented by M.Weng: www.mentaldev.com
- What about organizing a DevRob challenge on a simulator using PYRO? Details: Organizers provide a simulator, but without explaining what is simulated: this

simulator just gives, say 10 real values (sensors inputs) at each time step and collect 5 real values. The 5 ones are motor commands. But the participants don't know what kind of world is simulated (2D, 3D, 67D…), nor the kind of sensors they have, nor the kind of actuators they use. So they have to design a cognizer that will discover that thought simulation. Now how to decide who wins? (hehehe…)

General ideas about Developmental robotics

The only known system able to tackle **muddy** tasks (vision, speech, driving, real world issues) is the brain. **Biological inspiration** is a clue.
Everybody agree with the idea that a cognizer needs **learning** to face real world problems.
« Common sense knowledge must be grounded on **sensorimotor interaction** »
« Never start from scratch, but build on the old »
Almost all presented works are based on reinforcement learning.
"**Open-ended learning**": on line and forever learning.
"**Task general approach**": We don't specify a task to solve: the robot learns whatever it wants; build its own representation of its environment and of itself. And, finally, we hope that his learned world model will be powerful enough to solve muddy problems. At least it's funny.
We must **avoid to reify** and to transpose our model of the world on the robot representation. This idea is close to Brooks' approach. Indeed, as far as we don't have the same sensorimotor experience, we will never produce the same representation of the world. Even Time and Space notions are human! So what are the common foundations of cognizers? Information theory, statistics, mathematics. A DevRob fundamentalist should avoid giving any prior human knowledge to the robot, including about itself, its sensors geometry, its goals… Isn't it too much? Or shouldn't we make a progression in difficulties, and remove prior knowledge step by step?
Before the development, the robot should have a **structure** enabling a powerful enough expression power ("a dog will never talk"). Its body must be adapted to its ecology, and it has to have a good developmental program.
Learning by imitation and social interactions haven't been discussed, however they seem to be crucial in the mental development of animals. Is it the next step?
Continuing the life metaphor, our developmental program comes (maybe) from evolution: we were selected to be able to learn **incrementally**.

Controversial Topics

- **Frontiers** of the cognitive agent (remember J. Droulez with his Bayesian agent in last BIBA workshop): what is intrinsic/extrinsic, inside/outside.
- Do space and Time notions should be learned? Where is time coming from for a robot?
- Is it neo-behaviorism? We have to be careful with ideological **terminology**. And with terminology in general.
- The **evaluation** of that field is a big issue: complexity, realness of emerged behaviors? Usability of metaphors for psychologists? Ability of general task performing?

Philosophical issues for Developmental Robotics

Evolving AI

General discussion about AI and real problems. Pentti Kanerva.
We need a new kind of algorithm for computing in high dimensional spaces of real problems, we need to discover the secrets of High-D vectors, and for that mathematics are required, geometry and algebra.

Claim 1:

We must understand the brain, because it's the only system succeeding with such high-D problems. It's the technological challenge of the next century. As an analogy, it would have been impossible to succeed in celestial mechanics without understanding gravity, impossible to protect us against illness without microbiology, and to understand heredity without genetics.

Claim 2

Developmental robotics is our best bet.

How cognizers come to know their world and what this implies for cognitive robots

Alexander Kovacs
We have to work at the raw level, and not at the phenomenological level, do not reify, leave the things implicit and let the robot build its own representations.
Do not fear to have a lot of sensory inputs, because it is the solution, it's needed to build a good representation of the environment.

Learning about the self and others through contingency

A robot learns to recognize itself in its video inputs: Kevin Gold
This has been achieved considering that the self is what is moving immediately after a motor order is given. So the robot passed the mirror test. It has been shown that some monkeys are extending their body mental image to the tool there are using.

Toward learning the binding affordances of objects: A behavior-grounded approach

Developmental approach to learning the manipulation of objects
Alexander Stoytchev
Goal: have robots autonomously use tools to over pass their physical limitations. For that the robot tries longer and longer sequences of action.

A developmental approach to grasping

Lorenzo Natal

Active learning through interaction allows:

- To relate different sensory modalities together
- Autonomously drives exploration
- Establish casual links

Human are optimized for adaptation, not for performance.
What we call “simple tasks” like grasping, has been learned with a lot of trials and efforts
For finding his hand in his video flow, the robot uses background difference and correlation with motor commands (neural net).

Intrinsic Motivation

Auto-supervised learning in the Bayesian programming framework

Pierre Dangauthier

Intrinsically motivated reinforcement learning: A promising framework for developmental robot learning

Andrew Stout,

The Playground environment: Task-independent development of a curious robot

Curiosity for Sony dog robot: Pierre-Yves Oudeyer
Open-ended development can be driven either by human reward, either by intrinsic motivation. Motivation could come from: curiosity, surprise, challenge, novelty... But it’s not easy to measure those things, and different measure (correlation, information, entropy…) can lead to different appreciation of the novelty of a situation.
The chosen solution (Same as Schmidhuber 1991) is pleasure of learning, that is to say reward=high decrease in error rate of predictions made on the world. If error rate is low but constant, the robot gets “bored” and then tries another new sensorimotor experience. If error rate is high and does not decrease, it means that this part of the environment is unpredictable, at least with the embedded learning algorithm. Then it’s frustrating and the robot looks for something else. See website for videos, it’s really impressive how the robot behave like a baby.

Panel Session

Why developmental robotics?

- Engineering answer: improving task resolution
- Life science: to provide a metaphor of life by increasing behavioral complexity.

How to evaluate results?

- Engineering: if we develop models of the world enabling, in the future to tackle a lot of different tasks.

- Life science: if we provide useful metaphors for psychologist, qualitative comparison with human development. In this case psychologists should enter in the loop.

But theses two approaches can be contradictory.

Tack-independent learning

Robot generates its own goals and asks dynamically
It continuously adapts as environment changes
Reinforcement learning (RL) is natural in this case. Its remains to find a good reward from:

- Extrinsic: a teacher (human)
- Intrinsic:
    - Resource driven: pain, food, energy, restlessness
    - Information driven: novelty, boredom, curiosity
- Combination of both

To robot should look for both opposites predictability and novelty, and it won't oscillate between them because it learns.
Difference between novelty (never seen) and surprise (unexpected)

Misc.

Building blocks are a necessity; they should be learned with their hierarchy.

There is 2 sort of complexity:

- Computational: can't think forever
- Experimental: can't live forever.

Emotions are a heuristic to survive.
Neurophysiology is incredibly complex; some people are glued for their all life studying base ganglia neurons? We need to know about neurophysiology and psychology, but not to be stuck on them.

We provide metaphor of development, and not models, because a model should make predictions on a phenomenon of the world, and because a model should be based on data.

Developmental Architectures

Perceptual memory and learning: Recognizing, categorizing, and relating

Stan Franklin

A framework for the development of robot behavior

Roderic Grupen

The possibilities (kinematics and dynamics) of the body lower the dimensionality of the search space (e.g. babies are short sighted, can't move a lot). This is certainly not pure hazard, and can bean an important propriety of incremental development.
In this work, the robot learns incrementally more complex behaviors, reusing simpler ones stored as "schemas". Learning is done on the schema space, which is drastically smaller.

Automatic language acquisition by an autonomous robot

Stephen Levinson,
There is no brain without body, no isolated cognitive functions, language is acquired thought interaction with environment, and sensorimotor experience is essential.
Mental processes are largely based on associative memory.
The language engine is primarily semantic, not syntaxic.

Bootstrapping Commonsense Knowledge Through Sensorimotor Experience

Bootstrap Learning of Foundational Representations

Learning its own sensor geometry. Patrick Beeson,
How to go from "pixel ontology" to "object ontology"? Reliable predictions are impossible at the pixel level, abstraction needed.
Objects are a concise reliable abstraction of the robot experience, so it's impossible to be hand designed but it must be learned.
Map learning with uninterpreted sensors (AIJ Pierce 97). Really impressive work because the robot founds itself that he has sensors located all around him, just looking at their mutual correlations and using PCA.

Psychologically-Inspired Models of Development

An embodied mechanism for autonomous action selection and dynamic motivation

Lee Mccauley
Constructivist learning (Piaget) Conscientiousness (Baars)
Learning of (context-action-result) triples from near zero knowledge
I heard The Bad Word in this talk: "consciousness", but I didn't get the associated ideas.

Towards a what-and-where model of infants' object representations

Matthew Schlesinger

Tasks and developmental robotics

An emergent framework for self-motivation in developmental robotics

Douglas Blank

Problem of evaluation of our results, if task independence. Measuring the complexity of the learned internal representation isn't a good idea because it can't be compared across models, and because complexity of representation isn't necessarily linked to complexity of behaviors. The only measure is behavior (neo-behaviorism?), and Turing-like tests needed.

Muddy tasks and the necessity of autonomous mental development (AMD)

Juyang Weng

"Muddiness" is qualifying how a problem is difficult to solve by a computer. Chess is a "clean" task but vision, language processing, speech are muddy tasks. Very muddy tasks are difficult to both humans and computers (creation of knew knowledge, scientific and artistic invention…). Traditional programming and AI techniques failed with muddy tasks. Muddiness is defined as the product of several factors like size of inputs, size of outputs, uncertainty, rawness of inputs, clarity of goal (see paper)… Weng believes that AMD isn't just another AI technique promised to failure, but that AMD is necessary for really muddy tasks. For instance DARPA Grand Challenge is muddy in a lot of ways (except goal), and is therefore an interesting challenge for the community.