The Motivation for Self-Learning Embodied Agents

Jan Lemeire, VUB

October 2024

1 The goal

"*To bring robots out of well-controlled environments (into the wild), they should be able to adapt and learn.*"

The goal of a *self-learning embodied agent* is to let him start with an empty brain - without prior knowledge, just with an architecture and some algorithms - and to learn:

- the effect of its actions on the world
- to interact with the world
- to control its actuators
- to control its environment
- to make effective action plans to achieve certain goals

2 Why apply self-learning on an agent?

Programming the behavior of a robot based on our scientific and mathematical understanding of the world has its successes. So, why self-learning:

- we don't have to provide the knowledge; it will be discovered autonomously if successful
- to finetune behavior (like calibration)
- to adapt to new, unseen, and uncontrolled environments (called open environments)
- to modify actuator control in case its functionality changed (due to wear or damage, for instance)
- to become versatile
- to understand ourselves, see Section 5

3 Challenges for the architecture and algorithm

The ultimate goal of the domain is:

- achieving the same learning performance as humans, who learn generic knowledge from just a few examples, and
- finding the least amount of prior knowledge; the **first principles** (architecture, algorithms, knowledge) which makes self-learning generically possible.

4 Opportunities of embodiment for Machine Learning

In most ML/AI settings, one has to extract knowledge from a given fixed dataset. An embodied agent, on the other hand, can explore and decide which parts of the state space it wants to sample (active learning). A robot can select the most informative samples by considering the ambiguities and conflicts in its predictive model and history. It can test hypotheses of the world. Embodiment thus allows agents to learn incrementally by integrating exploration, exploitation, and learning.

[Ronald Meester Marc Jacobs - De onttovering van AI] AI-ontwikkelaars gaan nog te vaak uit van een cartesiaans dualisme, waarbij denken en fysieke realiteit als twee afzonderlijke entiteiten worden gezien. De beroemde uitspraak 'Ik denk, dus ik ben' van Rene Descartes is een misleidend uitgangspunt om AI (en ons denken) te begrijpen. Wij mensen zijn fysieke, tactiele wezens. Onze lijfleijkheid speelt een cruciale rol in de manier waarop we de werkelijkheid ervaren, interpreteren en verstaan.

5 Scientific importance

The self-learning challenge opens up a lot of questions that relate to several domains throughout the sciences, of which the most important is what we can learn about ourselves, humans. Because we are the example of good practices in self-learning; human babies are born without any skills. They even don't know how to control their muscles or limbs. They just have the skill to learn... And become experts in understanding and controlling the world. Better than animals that are preprogrammed with instincts. Or at least humans are more versatile.

The field's tentacles are touching questions of:

- **adaptiveness**: can we achieve more through learning?
- the innate, **first principles** of learning. Learning from scratch prevents being biased toward a structure that implicitly contains knowledge.
- **developmental psychology**: how humans learn. Can we mimic it? What can we learn from it? Can we better understand the challenges of human learning? How humans learn is different than the typical brute force machine learning approaches.

- the **brain's architecture** (neurology): how knowledge (models) are stored in our brain (see 1000 brain theory).
- the development of **language** and concept learning. Self-learning leads to a bottom-up approach which contrasts with top-down approaches in which humans design and build methods to control robots, store knowledge, and let them reason. We engineer how knowledge is stored, organized, and, finally, used to control the environment. With self-learning, the challenge is to find out how an embodied agent can learn this bottom-up: based on interactions with its environment. In our approach, the agents will learn pieces of knowledge by experience. The structural and qualitative nature of the models enables the link with language bottom-up. The advantage is that the agent knows how to link the symbolics to the real world (grounding), apply it, and verify it(!).
- **'analysis by synthesis'**: embodied agents can test hypotheses from psychology and cognitive science, provide insight, and find design principles of biological systems. Construct it to understand it.
- **intelligence, awareness, philosophy of science**...

It would be fabulous if this research could contribute to some of these topics.

6 The proposed approach: contextual causal qualitative models.

"*Mimic the way human infants play, learn and develop.*"

We put forward an architecture that is fundamentally different than the currently-most popular probabilistic monolithic black box neural network approaches (cf active inference). This approach is meant to be more effective and to capture better the way that humans learn.

Basically, instead of trying to learn and use detailed quantitative, probabilistic models, our approach relies on deterministic qualitative models in which context is modeled explicitly.