Representational redescription: the next challenge?

Stephane Doncieux

June 4, 2015

Drawing inspiration from developmental psychology, it has been suggested to build cognitive architectures that allow robots to progressively acquire abstract representations [4]. Humans don't have a single optimal representation of the problems they solve. They can redescribe the information they have acquired in different formats [5]. It allows them to explore different representations and use multiple problem solving strategies, from low-level systematic search to abstract reasoning [2].

Representational redescription is the ability to change the way information is stored and manipulated, to make further treatments easier and more efficient. A representation is the description of some data in a given format. The lowest level possible for formats is the raw format of sensors and effectors. Some examples of high level representation can be drawn from artificial intelligence and machine learning communities: markov decision processes formalism, first order logic or neural networks. Changing the representation allows usage of different problem solving strategies. Adapted representations make computations easier by relying on a small set of relevant primitives instead of a big set of unstructured data.

Use a single representation or change representations over time?

Humans may use representational redescription because of physiological constraints. The genome contains twenty thousands genes to describe the whole body, including the brain with its hundred billions of neurons. Such a small number of genes may not be enough for a genetic transmission of sophisticated representations. Does it necessarily mean that robots should follow the same path? Human representational redescription may also be an advantage rewarded by evolutionary pressure because of the adaptation ability it has resulted in. Would it help robots to face open environments? This would undoubtedly be an interesting feature. In the following, we will consider the questions that it raises.

Where to start?

Sensorimotor data first need to be observed before they can be redescribed in a format that allows an agent to better understand what happened and eventually to reproduce it. Babies have grasping or sucking reflexes that allow them to start interacting with surrounding objects before they can perform more complex actions. Guerin et al. suggest using a similar set of innate sensorimotor schemas to bootstrap the process [4]. How to choose this set of primitive schemas and where to stop? If we, as roboticists, do know how to implement an efficient grasping behavior, why should we start with an inefficient grasping reflex? A sophisticated grasping behavior may allow the robot to rapidly and efficiently interact with objects, thus generating a lot of useful data to learn about them. Where should we then put the frontier between the schemas that are provided to the robot and the ones that should be discovered? Providing efficient behaviors is clearly a convenient way to bootstrap the process. Are there other alternatives?

Evolution shaped development, but could it be also involved in the representational redescription process?

Evolution has shaped, over millions of years, living beings and their development process. But beyond this first evo-devo relation, evolutionary mechanisms may also be at play *during* development and learning. The principles of variation and selection have contributed to the success of evolutionary computation because of their simplicity, robustness and versatility. They have been used in a robotics context for more than twenty years [1], and were notably able to generate non trivial behaviors with neural networks. They are also believed to be the primary mechanisms in development, both for learning motor schemas and for selecting problem solving strategies [4]. They could then have a significant role to play in the representational redescription process, in particular thanks to their ability to generate controllers relying on the most simple representations, i.e. sensorimotor data. Furthermore, this hypothesis may be biologically plausible, as evolutionary principles can be implemented along with neural mechanisms [3]. Evolution could then be involved in brain functions and thus in development and learning.

Should representation formats be given a priori or should it emerge from the developmental process?

Representational redescription requires the availability of the representation formats in which the redescription is expected to occur. A first possibility would be to provide the agent with different representation formats, like first order logic or markov decision processes formalism, for instance. Dedicated machine learning algorithms could extract them from a lower level representation, e.g. the sensorimotor flow. An alternative would be to use a versatile connectionist formalism and rely on deep learning algorithms to redescribe lower layers representations to more abstract ones. The first alternative is a somewhat top-down approach in which learning and decision algorithms are available from the very beginning. The developmental process "just" needs to represent sensorimotor data in the corresponding format for the system to exhibit high level cognitive abilities. The second is a bottom-up approach in which higher level representations emerge progressively and where the corresponding problem solving strategies will also need to emerge.

Does provided knowledge limit developmental abilities?

Providing knowledge allows one to take shortcuts in the developmental process: no need to discover what is provided and the corresponding developmental time is then saved. Providing sensorimotor schemas or representation formats constrains what the agent can do, what it will observe and what it will extract from these observations. If the agent is expected to face an open environment, isn't it a limit to its adaptive abilities? Are there conflicts between shortening developmental time and having an open-ended developmental process?

How to make a robot endowed with representational redescription transparent?

Giving a robot the ability to change its representations and problem solving strategies may make it difficult to understand for a human. A non-expert may have trouble predicting what the system will actually do and what it understands from its environment. Making such robots transparents may then be critical for them to be used in practice, in particular if they are to enter our everyday environment. How could it be achieved?

References

[1] S. Doncieux, N. Bredeche, J.-B. Mouret, and A.E. Eiben. Evolutionary robotics: what, why, and where to. *Frontiers in Robotics and AI*, 2, 2015.

[2] J.St.B.T. Evans. In two minds: dual-process accounts of reasoning. *Trends in Cognitive Sciences*, 7(10):454–459, October 2003.

[3] C. Fernando, E. Szathmary, and P. Husbands. Selectionist and evolutionary approaches to brain function: a critical appraisal. *Frontiers in Computational Neuroscience*, 6(April):1–28, 2012.

[4] F. Guerin, N. Kruger, and D. Kraft. A Survey of the Ontogeny of Tool Use : from Sensorimotor Experience to Planning. *IEEE Transactions on Autonomous Mental Development*, 5(1):18–45, 2013.

[5] A. Karmiloff-Smith. *Beyond Modularity: A Developmental Perspective on Cognitive Science.* The MIT Press, 1995.