Social Learning in DREAM

Julien Hubert and Jacqueline Heinerman and Evert Haasdijk
Vrije Universiteit Amsterdam, The Netherlands
{j.g.hubert|j.v.heinerman|e.haasdijk}@vu.nl

The DREAM project [1] develops a cognitive architecture that incorporates dream-like processes to allow robots to redescribe their knowledge and experiences at more abstract levels. Thus, the robots can generalise their experiences to learn and adapt more efficiently.

DREAM considers individual robots as well as collectives of robots. In a collective setting, the robots can combine forces to learn jointly, i.e., the robots implement social learning (SL). Social learning in the context of different levels of abstraction in the knowledge base poses an intriguing question: under which circumstances is it better for the robots to exchange knowledge at the abstract or at the immediate level?

Consider, for instance, a task in which a group of robots must explore a maze. This particular maze possesses cues indicating which direction to follow to reach the exit, e.g. a blue mark on a wall indicates turning left at the next crossing. A single robot could navigate the maze and learn the association between the cues and the path to follow, but that would require it to explore the complete maze. A group of robots could independently learn the correlations present in the environment and share their knowledge to increase their performance in solving the maze.

To accomplish any task in SL, the robots must exchange information, and one important question to answer is where this information originates. Should it be taken from the low level of the controller, e.g. the sensors, or from a higher level of abstraction. Consider that the controller of the robot is an artificial neural network (ANN). In that case, the weights of the ANN can be transferred to the other robots. But if they do not share the exact same sensors, there could be confusion in the integration of the new information, which would lead to a lower performance. Instead, we could use a higher level of abstraction, such as "if blue, turn right", but, similarly, the notion of blueness might differ between robots. These examples show that there is no clear answer to our question.

At the sensorimotor level, it would be impractical to stream the entire perception of a robot to its group, and let all robots learn from it. A better approach could be to decompose the stream into multiple experiences, and share these with other robots requiring them. How to decompose experiences is an open question in SL that we started exploring. Our approach is to represent each experience by a data point, and cluster them in sensory space. The set of points in a cluster form a distribution representing a *situation*. Each new experience can be compared to already detected situations and, eventually, be recognised if it has already been encountered. If not, a new cluster is created. From this library of situations, it becomes possible to exchange experiences with other robots. If the sensory information is similar, a robot can recognize a situation experienced by another and ask for additional information, e.g. the skill to act upon it. If the sensors differ, the recipient will dismiss the transmission.

From this framework, it becomes possible to exchange sensory experiences, but it also offers the possibility to build a higher abstraction level where these would be rewritten in a symbolic form. Ultimately, sharing the sensory experience and their symbolic equivalent could allow robots to create a sensory mapping with other robots, and reduce the problem caused by the difference in symbol grounding.

[1] Funded by European Union FET Proactive Initiative: Knowing, Doing, Being: Cognition Beyond Problem Solving under grant agreement 640891; http://www.robotsthatdream.eu