

Interactive Learning of Grounded Concepts

Jens Nevens, Paul Van Eecke, and Katrien Beuls

Artificial Intelligence Lab
Vrije Universiteit Brussel
Pleinlaan 2, B-1050 Brussels, Belgium
{jens|paul|katrien}@ai.vub.ac.be

Autonomous agents perceive the world through streams of continuous sensori-motor data. Yet, in order to reason and communicate about their environment, agents need to be able to distil meaningful symbolic concepts from their raw observations. Without such a repertoire of concepts, communication would need to happen by directly sending sensori-motor values. Such a system easily leads to miscommunication when perfect calibration is not possible.

Existing approaches to bridge between the continuous and symbolic domain include deep learning techniques (e.g. [1]) and version space learning [2]. Deep learning techniques generally achieve high levels of accuracy. However, they rely on very large amounts of training data, they often fail to adapt to unseen scenarios, and the resulting concepts lack transparency. Version space learning, on the other hand, can yield human-interpretable concept representations, but are notoriously brittle when faced with noisy training data.

In this interactive demo, we introduce a novel approach to grounded concept learning. Using the language game methodology [3], we set up a tutor-learner scenario where the learner is an autonomous agent, grounded in the world using a Nao humanoid robot, and the participant is its tutor. Using blocks of various shapes, sizes and colours, the tutor first creates a scene. In this scene, the tutor chooses a topic and describes it to the learner using an informative concept, such as ‘red’ or ‘cube’. The learner robot observes the world through human-interpretable streams of numeric data, such as ‘area’, ‘colour’ and ‘XY-coordinates’. These are obtained through standard computer vision techniques. The robot tries to find the object described by the tutor. After each interaction, whether is was successful or not, the tutor provides feedback to the learner by showing the intended topic object. The robot maximally benefits from this feedback to newly create or extend the representation of the concept. For each concept, the robot has to find out which data streams are important and what the typical values for each data stream within a concept are. To make these decisions, the learner makes use of the notion of *discrimination*, i.e. separating one particular object from the other objects in the scene. Over the course of many such interactions, the learner incrementally and in real-time builds a complete repertoire of concepts that is functional in the world. An overview of the experimental set-up is shown in Figure 1. A video of the demonstration can be found at <https://ehai.ai.vub.ac.be/demos/interactive-concept-learning>.

Copyright © 2019 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0)




Fig. 1. Overview of the experimental setup. The Nao humanoid robot observes a scene of blocks with different shapes, colours and sizes.

During the demonstration, participants are able to inspect the conceptual system of the agent and follow its evolution. The acquired concepts are completely transparent and human-interpretable. We show that our approach does not rely on huge amounts of training data, since forming a repertoire of concepts only requires a few interactions. Additionally, the resulting concepts are general enough to be applied to previously unseen objects and can be learned in an incremental manner. The whole system is adaptive as it does not require us to specify the number of concepts that should be learned. This completely depends on the objects observed by the agent, hence there is no need for complete or even partial retraining when the environment changes. These properties make the approach well-suited to be used in robotic agents as the module that maps from continuous sensori-motor input to grounded, symbolic concepts that can then be used for higher-level reasoning tasks such as planning, explanation or communication.

References

1. Higgins, I., Matthey, L., Glorot, X., Pal, A., Uria, B., Blundell, C., Mohamed, S., Lerchner, A.: Early visual concept learning with unsupervised deep learning. arXiv preprint arXiv:1606.05579 (2016)
2. Mitchell, T.M.: Generalization as search. *Artificial intelligence* **18**(2), 203–226 (1982)
3. Steels, L.: Language games for autonomous robots. *IEEE Intelligent systems* **16**(5), 16–22 (2001)