



Editorial: Robots that Learn and Reason: Towards Learning Logic Rules from Noisy Data

Plinio Moreno^{1*}, Alexandre Bernardino¹, José Santos-Victor¹, Rodrigo Ventura¹ and Kristian Kersting²

¹Institute for Systems and Robotics, Instituto Superior Técnico, Universidade de Lisboa, Lisbon, Portugal, ²Centre for Cognitive Science, Computer Science Department, TU Darmstadt, Darmstadt, Germany

Keywords: learning logic rules, robotics, predicate grounding, two-level planning, reinforcement learning

Editorial on the Research Topic

Robots that Learn and Reason: Towards Learning Logic Rules from Noisy Data

From the early developments of AI applied to robotics by Hart et al. (1968), Duda and Hart (1972) and Lozano-Pérez and Wesley (1979), higher level commands were grounded to real world sensing by carefully design algorithms, which provide a link between the abstract predicates and the sensors and actuators. In order to have fully autonomous robots that learn by exploration and by imitation, the grounding algorithms between the higher-level predicates and the lower-level sensors and actuators should be discovered by the robot. Previous and recent efforts on robotics aim to discover and/or learn these intermediate layer commands, which must cope with discrete and continuous data. The main objective of this Research Topic is to advance on learning logic rules from noisy data. We have four articles that address: Logic rules that cope with states that are not directly observable by the sensing modalities; learning rules that represent object properties and their functionalities, which are grounded to the particular robot experience; learning low-level robot control actions that fulfill a set of abstract predicates in a two-level planning approach; learning to develop skills in a robotic playing scenario by composing a set of behaviors. In the following, we introduce the four articles and their contributions to rule learning in presence of noisy data.

The article by Zuidberg Dos Martires et al. introduces a general-purpose semantic object tracker, which anchors sensory input to object representations in a probabilistic manner. The semantic tracker is able to estimate the object position in the presence of total occlusions of the object by using a multi-modal probability distribution, which is modeled as a distributional logic tree. In addition, the tree structure of the multi-modal distribution is estimated from observations, which provides a set of data-driven clauses that can be augmented with recursive ones. This methodology allows to model objects that can be occluded by objects that are themselves occluded.

The article by Thosar et al. defines measurable object properties and their functionalities, focusing on the robot-centered estimation of the object properties for multiple categories. Regarding the object properties, the robot executes actions that interact with the objects, measuring the effects on the objects through proprioceptive changes in the robot and external sensors. The robot-centric conceptual knowledge is obtained from unsupervised clustering of the numerical values, where each cluster represents a qualitative property (i.e., concept). Regarding the object functionalities, these are selected according to a set of predefined tasks, where a quantitative value of the functionality is obtained from vision-based sensors. The object properties and their functionalities are designed for selecting a tool substitute, which is addressed as a classification problem. The learning algorithm selects tool substitutes very similar to human experts, showing the validity of the approach. This methodology allows to select tool substitutes in a robot-centric manner, based on unsupervised discovery of qualitative object properties.

OPEN ACCESS

Edited and reviewed by:

Mikhail Prokopenko,
The University of Sydney, Australia

*Correspondence:

Plinio Moreno
plinio@isr.tecnico.ulisboa.pt

Specialty section:

This article was submitted to
Computational Intelligence in Robotics,
a section of the journal
Frontiers in Robotics and AI

Received: 09 August 2021

Accepted: 12 August 2021

Published: 16 September 2021

Citation:

Moreno P, Bernardino A,
Santos-Victor J, Ventura R and
Kersting K (2021) Editorial: Robots that
Learn and Reason: Towards Learning
Logic Rules from Noisy Data.
Front. Robot. AI 8:755933.
doi: 10.3389/frobt.2021.755933

The article by Eppe et al. proposes a methodology that integrates an abstract (high-level) planner using Planning Domain Definition Language (PDDL), and an actor-critic Deep Neural Network (low-level) using Hindsight Experience Replay (HER). On one hand, model-free Reinforcement Learning (RL) with Universal Value Function Approximators (UVFAs) and HER provide the tools to build a model that grounds a discrete symbolic action space in continuous action trajectories. The hybrid nature of the approach allows to cope with complex casual dependencies, such as stacking three blocks in experimental conditions with realistic noise. On the other hand, although the predicates that correspond to the discrete symbolic space are handcrafted properly for the scenarios, the low-level subgoal g_{sub} works across predicates. This design allows to have only one RL model that executes actions for all the created predicates. This hybrid two-level planner works better than the more general Hierarchical Reinforcement Learning (HRL) for the ant navigation scenario, but the application to other domains requires always the handcrafted predicate design.

The article by Hangl et al. proposes a methodology that improves previously developed work on skill learning, which relies on composition of a fixed set of simple behaviors. The main contribution of this article is to discover the sequence (i.e., composition) of simple behaviors that leads to the goal, which in previous work was set by an expert. The behaviors are learnt in a bottom-up hierarchical fashion by running the sensing \rightarrow preparation \rightarrow basic behavior execution \rightarrow reward loops, using the Projective Simulation (PS) framework that includes: 1) actions

for sensing, 2) perceptual states and 3) preparatory behaviors. The basic building block of PS is a clip, which creates connections between the perceptual states and the behaviors. On top of the clips, a skill is defined as a behavior or a composition of behaviors that lead to successful execution. To acquire new skills, a random walk through the PS framework is performed by hopping from clip to clip until a behavior is reached. The robot learns by adapting the probabilities of a clip given the previously executed clip, according to the received reward. Once a behavior is reached with a high accuracy, a new behavior is explored as continuation of the previous one. Complex manipulation behaviors are discovered by the model, such as book grasping and tower disassembling. However, the behavior compositions are limited to perceptual states of the same sensing behavior.

This research topic contributes to the very challenging problem of understanding the world from perceptual input while acquiring predicates that describe the task-environment relations. The selected articles provide insights on rule learning for: 1) Non-observable states during visual tracking, 2) representation of objects and their properties in a robot-embodied way, 3) developing skills in a robotic playing scenario, and 4) low-level robot control actions that fulfill predicates in a two-level planning approach.

AUTHOR CONTRIBUTIONS

All authors listed have made a substantial, direct, and intellectual contribution to the work and approved it for publication.

REFERENCES

- Duda, R. O., and Hart, P. E. (1972). Use of the Hough Transformation to Detect Lines and Curves in Pictures. *Commun. ACM* 15, 11–15. doi:10.1145/361237.361242
- Hart, P., Nilsson, N., and Raphael, B. (1968). A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Trans. Syst. Sci. Cyber.* 4, 100–107. doi:10.1109/TSSC.1968.300136
- Lozano-Pérez, T., and Wesley, M. A. (1979). An Algorithm for Planning Collision-free Paths Among Polyhedral Obstacles. *Commun. ACM* 22, 560–570. doi:10.1145/359156.359164

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2021 Moreno, Bernardino, Santos-Victor, Ventura and Kersting. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.