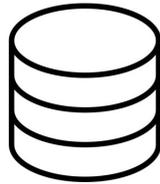


A method for self-supervised robotic learning that entails setting feasible goals

15 November 2019, by Ingrid Fadelli



$$\{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_N\}$$

The robot collects random interaction data to be used for training a representation and as off-policy data for RL. Credit: Nair et al.

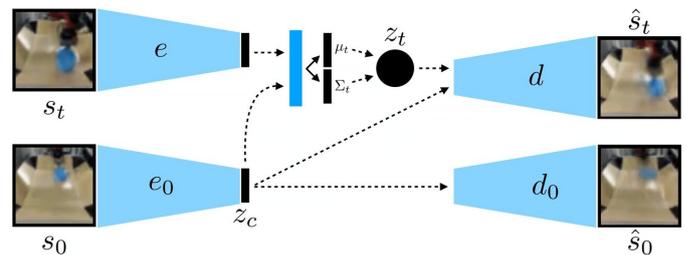
Reinforcement learning (RL) has so far proved to be an effective technique for training artificial agents on individual tasks. However, when it comes to training multi-purpose robots, which should be able to complete a variety of tasks that require different skills, most existing RL approaches are far from ideal.

With this in mind, a team of researchers at UC Berkeley has recently developed a new RL approach that could be used to teach robots to adapt their behavior based on the task with which they are presented. This approach, outlined in [a paper pre-published on arXiv](#) and presented at this year's Conference on Robot Learning, allows robots to automatically come up with behaviors and practice them over time, learning which ones can be performed in a given environment. The robots can then repurpose the knowledge they acquired and apply it to new tasks that human users ask them to complete.

"We are convinced that data is key for robotic manipulation and to obtain enough data to solve manipulation in a general way, robots will have to collect data themselves," Ashvin Nair, one of the

researchers who carried out the study, told TechXplore. "This is what we call self-supervised [robot learning](#): A [robot](#) that can actively collect coherent exploration data and understand on its own whether it has succeeded or failed at tasks in order to learn new skills."

The new approach developed by Nair and his colleagues is based on a goal-conditioned RL framework [presented in their previous work](#). In this previous study, the researchers introduced goal setting in a latent space as a technique to train robots on skills such as pushing objects or opening doors directly from pixels, without the need for an external reward function or state estimation.



The researchers trained a context-conditioned VAE on the data, which disentangles context that stays constant during a rollout. Credit: Nair et al.

"In our new work, we focus on generalization: How can we do self-supervised learning to not only learn a single skill, but also be able to generalize to visual diversity while performing that skill?" Nair said. "We believe that the ability to generalize to new situations will be key for better robotic manipulation."

Rather than training a robot on many skills individually, the conditional goal-setting model proposed by Nair and his colleagues is designed to set specific goals that are feasible for the robot and

are aligned with its current state. Essentially, the algorithm they developed learns a specific type of representation that separates things that the robot can control from the things it cannot control.

When using their self-supervised learning method, the robot initially collects data (i.e. a set of images and actions) by randomly interacting with its surrounding environment. Subsequently, it trains a compressed representation of this data that converts images into low-dimensional vectors that implicitly contain information such as the position of objects. Rather than being explicitly told what to learn, this representation automatically understands concepts via its compression objective.

"Using the learned representation, the robot practices reaching different goals and trains a policy using [reinforcement learning](#)," Nair explained. "The compressed representation is key for this practice phase: it is used to measure how close two images are so the robot knows when it has succeeded or failed, and it is used to sample goals for the robot to practice. At test time, it can then match a goal image specified by a human by executing its learned policy."

The researchers evaluated the effectiveness of their approach in a series of experiments in which an artificial agent manipulated previously unseen objects in an environment created using the MuJoCo simulation platform. Interestingly, their training method allowed the robotic agent to acquire skills automatically that it could then apply to new situations. More specifically, the robot was able to manipulate a variety of objects, generalizing manipulation strategies it previously acquired to new objects that it had not encountered during training.

"We are most excited about two results from this work," Nair said. "First, we found that we can train a policy to push objects in the real world on about 20 objects, but the learned policy can actually push other objects too. This type of generalization is the main promise of deep learning methods, and we hope this is the start of much more impressive forms of generalization to come."

Remarkably, in their experiments, Nair and his

colleagues were able to train a policy from a fixed dataset of interactions without having to collect a large amount of data online. This is an important achievement, as data collection for robotics research is generally very expensive, and being able to learn skills from fixed datasets makes their approach far more practical.

In the future, the model for self-supervised learning developed by the researchers could aid the development of robots that can tackle a wider variety of tasks without training on a large set of skills individually. In the meantime, Nair and his colleagues plan to continue testing their approach in simulated environments, while also investigating ways in which it could be enhanced further.

"We are now pursuing a few different lines of research, including solving tasks with a much greater amount of visual diversity, as well as solving a large set of tasks simultaneously and seeing if we are able to use the solution on one task to speed up solving the next task," Nair said.

More information: Contextual imagined goals for self-supervised robotic learning. arXiv:1910.11670 [cs.RO]. arxiv.org/abs/1910.11670

Visual reinforcement learning with imagined goals. arXiv:1807.04742 [cs.LG]. arxiv.org/abs/1807.04742

ccrig.github.io/

© 2019 Science X Network

APA citation: A method for self-supervised robotic learning that entails setting feasible goals (2019, November 15) retrieved 1 July 2022 from <https://techxplore.com/news/2019-11-method-self-supervised-robotic-entails-feasible.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.