

An algorithm to teach robots pre-grasping manipulation strategies

August 12 2019, by Ingrid Fadelli



Credit: Berscheid, Meissner & Kröger.

When human beings reach out to grasp a given object, they often need to push clutter out of the way in order to isolate it and ensure that there is enough room to pick it up. Even though humans are not always fully

aware that they are doing it, this strategy, known as "pre-grasping manipulation," allows them to grasp objects more efficiently.

In recent years, several researchers have tried to reproduce human manipulation strategies in robots, yet fewer studies have focused on pre-grasping manipulation. With this in mind, a team of researchers at Karlsruhe Institute of Technology (KIT) has recently developed an algorithm that can be used to train robots on both grasping and pre-grasping manipulation strategies. This new approach was presented in a paper [pre-published on arXiv](#).

"While grasping is a well-understood task in robotics, targeted pre-grasping manipulation is still very challenging," Lars Berscheid, one of the researchers who carried out the study, told TechXplore. "This makes it very hard for robots to grasp objects out of clutter or tight spaces at the moment. However, with the recent innovations in machine and [robot learning](#), robots can learn how to solve various tasks by interacting with its environment. In this study, we wanted to apply an approach we presented in our [prior work](#) not only to grasping, but to pre-grasping manipulation as well."



Credit: Berscheid, Meissner & Kröger.

When a robot is learning how to complete a certain task, it essentially needs to figure out how to solve a problem by maximizing its rewards. In their study, the researchers focused on a task that involved grasping objects out of a randomly filled bin.

The robot was trained on how to grasp objects for approximately 80 hours, using input from a camera and feedback from its gripper. When it successfully held an object in its robotic gripper, it attained a reward. The algorithm developed by Berscheid and his colleagues takes the robot's training one step further, allowing it to also acquire useful for pre-grasping manipulation strategies, such as shifting or pushing.

"The key idea of our work was to extend the grasping actions by introducing additional shifting or pushing motions," Berscheid explained. "The robot can then decide what action to apply in different situations. Training robots in reality is very tricky: First, it takes a long time, so the training itself needs to be automated and self-supervised, and second a lot of unexpected things can happen if the robot explores its environment. Similar to other techniques in machine learning, robot learning is always limited by its data consumption. In other words, our work is connected to two very challenging research questions: How can a robot learn as fast as possible—and what tasks can a robot learn using the discovered insights?"



Credit: Berscheid, Meissner & Kröger.

As Berscheid goes on to explain, a robot can learn more efficiently if it receives direct feedback after each action it performs, as this overcomes the issue of sparse rewards. In other words, the more feedback provided to a robot (i.e. the more rewards it receives for successful actions), the faster and more effectively it learns how to complete a given task.

"This sound easy, but is sometimes tricky to implement: For example, how do you define the quality of a pre-grasping manipulation?" Berscheid said.

The approach proposed by the researchers is based on a previous study

that investigated the use of differences in grasping probabilities before and after a particular action, focusing on a small area around where the action is performed. In their new study, Berscheid and his colleagues also tried to uncover actions that a robot should try to learn as fast as possible.

"This is the well-known problem of exploration in robot learning," Berscheid explained. "We define an exploration strategy that either maximizes the self-information or minimizes the uncertainty of actions and can be computed very efficiently."

The algorithm presented by the researchers allows a robot to learn the optimal pose for pre-grasping actions such as clamping or shifting, as well as how to perform these actions to increase the probability of successful grasping. Their approach makes one particular action (i.e. shifting) dependent on the other (i.e. grasping), which ultimately removes the need for sparse rewards and enables more efficient learning.

The researchers applied their algorithm to a [Franka robotic arm](#) and then evaluated its performance on a task that involves picking up objects from a bin until it is completely empty. They trained the system using 25,000 different grasp and 2,500 shift actions. Their findings were very promising, with the robotic arm successfully grasping and filing both objects it was familiar with and others that it had never encountered before.

"I find two results of our work to be in particular exciting," Berscheid said. "First, we think that this work really shows the capability of robot learning. Instead of programming how to do something, we tell the robot what to do—and it needs to figure out how to do it by itself. In this regard, we were able to apply and generalize the methods that we've developed for grasping towards pre-grasping manipulation. Second and of more practical relevance, this could be very useful in the automation

of many industrial tasks, particularly for bin picking, where the robot should be able to empty the bin completely on its own."

In the future, the approach developed by Berscheid and his colleagues could be applied to other robotic platforms, enhancing their pre-grasping and grasping manipulation skills. The researchers are now planning to carry out further studies exploring other research questions.

For instance, so far their approach only allows the Frank robotic arm to grasp objects with an upright hand, using what are referred to as 'planar grasps'. The researchers would like to extend their algorithm to also enable lateral grasps, by introducing more parameters and using additional training data. According to Berscheid, the main challenge when trying to achieve this will be ensuring that the robot acquires lateral grasps while keeping the number of grasp attempts it performs constant during the training phase.

"In addition, grasping objects is often part of a high-level task, e.g. we want to place the object at a specific position," Berscheid said. "How can we place an unknown object precisely? I think that the answer to this question is very important to tackle both industrial and new applications in service robotics. In our project we want to keep the focus on real-world [robot](#) learning, bridging the gap between toy-examples in research and complex real-world applications."

More information: Robot learning of shifting objects for grasping in cluttered environments. arXiv: 1907.11035 [cs.RO].

arxiv.org/abs/1907.11035

Improving data efficiency of self-supervised learning for robotic grasping. arXiv:1903.00228 [cs.RO]. arxiv.org/abs/1903.00228

© 2019 Science X Network

Citation: An algorithm to teach robots pre-grasping manipulation strategies (2019, August 12)
retrieved 25 April 2024 from

<https://techxplore.com/news/2019-08-algorithm-robots-pre-grasping-strategies.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.