REINFORCEMENT LEARNING ALGORITHMS FOR CONTACT-RICH MANIPULATION TASKS

Project description

The automatization of manufacturing processes, such as assembly tasks, is pivotal in industrial applications. Automatizing assembly tasks is particularly complex due to two main reasons. Firstly, assembly tasks involve the so-called contact-rich manipulation operations, such as insertion and screwing, during which there are contacts between the components to be assembled. Secondly, position and dimension of the components are known with uncertainty, typically higher than the clearance required to accomplish tasks. Humans handle these two problems by exploiting information elaborated by our perception system to adapt to the external environment. In contrast, contact-rich manipulation tasks can be particularly complex for robots, which are not equipped with humans' remarkable perception and compliance capabilities.

Consequently, the traditional solution implemented to accomplish assembly tasks is to realize a special-purpose machine [1], a one-of-a-kind machine designed to assemble a specific product. In these machines, properly designed mechanisms place each object, at the right time, in a fixed position to compensate for uncertainties. This strategy is opposite to the approach followed by humans, based on adapting to the external environment. The drawbacks of special-purpose machines are the costs of realization, installation, maintenance, and adaptations to each product modification. Consequently, these machines are economically feasible when the product life is long, production volume is high, and there are minimal product variations over time.

The employment of robotic manipulators for contact-rich manipulation tasks is an interesting alternative to special-purpose machines. Robotic manipulators are much more flexible than special-purpose machines since they can be re-programmed in case of product modifications or at the end of the product life. However, accomplishing assembly tasks with a robotic manipulator in the presence of uncertainty and tolerances is particularly challenging. Over the years, several benchmarks have been proposed, see, for instance, the benchmark proposed by Siemens (Figure 1). Despite recent advances, the use of manipulators in unstructured environments, like in assembly tasks, is still a challenging problem worth investigating.

The high-level strategy implemented with robotic manipulators consists of equipping robots with a sufficiently informative sensing system and designing a reactive control strategy that, as performed by humans, modifies the robot's motion in accordance [1,2,3,4]. However, the derivation of reactive controllers by hand is complex and time-consuming. Recent advances have shown Reinforcement Learning (RL) potentialities in the automatic resolution of several complex control problems. RL algorithms learn to accomplish a task by optimizing a cost function based on outcomes collected while interacting with the environment [5].

Most RL-based solutions proposed for contact-rich manipulation resort to the so-called Model-Free RL (MFRL) [6,7,8]. Conveniently, MFRL algorithms do not derive a model of the system dynamics, which, in these setups, is particularly complex. Despite interesting results, the poor data efficiency of MFRL, i.e., the high number of tests on the real system required to converge, limits its applicability.

Model-Based RL (MBRL) is a data-efficient alternative to MFRL [9]. MBRL exploits collected data to build a system evolution model and optimizes the policy on the model instead of the actual system.

Deriving a model that accurately describes the complex dynamics of the system is the most challenging aspect of MBRL for contact-rich manipulation tasks. For this reason, there are only a few examples of MBRL algorithms for contact-rich manipulation tasks [10].

The ultimate goal of this project is deriving novel MBRL solutions to solve contact-rich manipulation tasks, with particular attention to two modeling aspects neglected by the few MBRL algorithms proposed in this context.

- The first one is the derivation of grey-box models of contact-rich manipulation systems merging information from data with prior knowledge of the system.
- The second one is the inclusion of physical constraints in the derived models. For instance, consider the peg-in-hole task. The peg motion is constrained by contacts with the hole and the environment. This project aims at deriving models that account for this kind of constraints when they simulate the system for policy optimization, thus avoiding that the policy plans unfeasible control strategies.

By leveraging grey-box models and inclusion of physical constraints, the proposed approach aims to improve modeling performance and speed up the whole learning process. As regards the methodology, we envision to rely on Bayesian approaches for the modeling tasks, in particular Gaussian Process Regression [11], which already proved remarkable data efficiency in robotic applications, see, for instance [12] in actual RL applications.

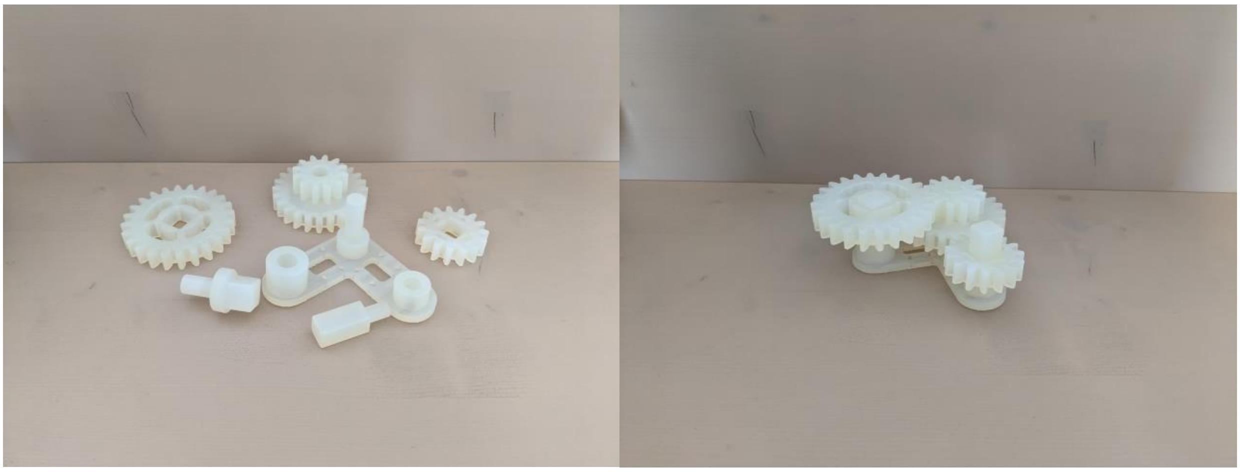

FIGURE 1: Pictures of the Siemens benchmark before and after assembling.

References

[1] J.L. Nevins and D.E. Whitney. Research on advanced assembly automation. Computer, 1977

[2] H. Bruyninckx, S. Dutre, and J. De Schutter. Peg-on-hole: a model based solution to peg and hole alignment. In Proceedings of 1995 IEEE Int. Conf. on Robotics and Automation, 1995.

[3] Wyatt S. N., Yonghong Z., and Yoh-Han P. Interpretation of force and moment signals for compliant peg-in-hole assembly. In Proceedings of the 2001 IEEE Int. Conf. on Robotics and Automation, ICRA 2001, May 21-26, 2001, Seoul, Korea. IEEE, 2001.

[4] D. E. Whitney. Quasi-Static Assembly of Compliantly Supported Rigid Parts. Journal of Dynamic Systems, Measurement, and Control, 03 1982

[5] D. P Bertsekas. Reinforcement learning and optimal control. Athena Scientific optimization and computation series. Athena Scientific, Belmont, Massachusetts, 2019

[6] T. Inoue, G. De Magistris, A. Munawar, T. Yokoya, and R. Tachibana. Deep reinforcement learning for high precision assembly tasks. In 2017 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), 2017

[7] J. Luo, E. Solowjow, C. Wen, J. Ojea, and A.M. Agogino. Deep reinforcement learning for robotic assembly of mixed deformable and rigid objects. In 2018 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), 2018

[8] G. Schoettler, A. Nair, J. Luo, S. Bahl, J. Aparicio Ojea, E. Solowjow, and S. Levine. Deep reinforcement learning for industrial insertion tasks with visual inputs and natural rewards. In 2020 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), 2020.

[9] K. Chatzilygeroudis, V. Vassiliades, F. Stulp, S. Calinon, and J. Mouret. A survey on policy search algorithms for learning robot controllers in a handful of trials. IEEE Trans. on Robotics, 2020.

[10] S. Levine, N. Wagener, and P. Abbeel. Learning contact-rich manipulation skills with guided policy search. In 2015 IEEE Int. Conf. on Robotics and Automation (ICRA), 2015

[11] C. E. Rasmussen and C. K. I. Williams. Gaussian Processes for Machine Learning. The MIT Press, 11 2005

[12] F. Amadio, A. Dalla Libera, R. Antonello, D. Nikovski, R: Carli, and D. Romeres. Model-based policy search using monte arlo gradient estimation with real systems application. IEEE Trans. on Robotics, 2022.