Deep Learning for Manipulation with Visual and Haptic Feedback

Sabrina Hoppe^{1,2}, Zhongyu Lou¹, Daniel Hennes², Marc Toussaint²

Abstract-Recent advances in deep learning for robotics have demonstrated the possibility to learn a mapping from raw visual input to control signals. For contact-rich real-world manipulation tasks however, it is questionable whether purely vision-guided control is sufficient. Aiming at a deep learning framework for deep imitation or reinforcement learning for manipulation from both visual and haptic feedback, we have investigated a peg-in-hole task with sensory feedback from a camera and a module providing both passive compliance and sensor feedback about the end effector displacement. We have trained a neural network that adjusts the end effector position on a horizontal plane while the height of the end effector is steadily decreased by a simple external controller. Preliminary results demonstrate that network performance increases when tactile feedback is available but leave several questions open for discussion and future investigations.

I. INTRODUCTION AND RELATED WORK

Recent work in deep reinforcement or imitation learning has demonstrated the possibility to train policies end-to-end from raw images to control signals directly [1]. However, it remains unclear whether policies as a function of visual input scale to more complex contact-rich manipulations. In particular, optimal controllers for high-dimensional contact-seeking behavior might be unknown. Therefore, most supervised learning approaches are infeasible for these cases. In analogy to human sensing behavior, it seems natural though to expect haptics to play a crucial role in (learning) manipulation tasks. While there is a large body of work explicitly modelling object contacts for manipulation [2], only few have integrated multimodal feedback into end-to-end learning systems for manipulation [3]. Tactile sensing has also been indirectly incorporated through force control [4].

As a stepping stone towards deep reinforcement learning from both visual and haptic feedback, we here present initial results on deep learning for an exemplary peg in hole task with passive compliance using both end effector displacement as well as camera images as feedback.

II. SYSTEM OVERVIEW

We are using a dual arm KAWADA Nextage robot (see Figure 1) for all experiments. A custom passive compliance module is mounted on one of the wrists. It provides full 6D feedback about the end effector's current positional and rotational displacement. The second arm provides the light source and camera images from a static view point for our experiments.

¹Robert Bosch GmbH, Robert Bosch Campus 1, Renningen, Germany. firstname.lastname@de.bosch.com





Fig. 1. System overview (left) and network architecture (right).

III. METHOD

Assuming that the object position is known during training, we have uniformly sampled 10,000 positions from a fixed space above the target object as well as 5,000 positions with surface contact. The real robot was moved to each position to collect an image as well as the compliance element's feedback at this point. For the contact samples, an initial position from where to start moving to the points was randomly chosen in order to enforce different contact angles. Based on these images, we have trained a neural network as shown in Figure 1 to map from an image and compliance feedback to the offset in x and y direction from the center of the target object. At test time, the height of the end effector was automatically decreased by an open-loop controller.

The first convolutional layer of the neural network is initialised with weights from GoogleLeNet trained on ImageNet [5]. The part of our network that depends on visual input only (black modules in Figure 1) serves both as a baseline and pretraining step (dashed arrow). To incorporate haptic feedback, the network is extended (gray modules) and the additional parameters are trained while the pretrained vision-based part of the network is frozen.

IV. RESULTS & OPEN QUESTIONS

For the network using camera input only, 77 out of 100 trials were successful. Using a closed-loop controller that goes up whenever the end effector got stuck, the success rate increases to 85%. The full network architecture trained on visual and haptic feedback solves the task in all 100 trails.

Open questions for future investigations include (1) generalisation of this approach, e.g. to new target positions and rotations; (2) generalisation to more complex, contactricher manipulation tasks for which we expect leveraging information from negative samples to be crucial; and (3) design choices about whether or not, and how to integrate the height controller into the network policy.

ACKNOWLEDGMENT

The authors thank Oleksandra Pariy who helped to set up the system and Hung Ngo for valuable discussions.

References

- [1] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1334–1373, 2016.
- [2] J. Tegin and J. Wikander, "Tactile sensing in intelligent robotic manipulation-a review," *Industrial Robot: An International Journal*, vol. 32, no. 1, pp. 64–70, 2005.
- [3] H. van Hoof, N. Chen, M. Karl, P. van der Smagt, and J. Peters, "Stable reinforcement learning with autoencoders for tactile and visual data," in *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*. IEEE, 2016, pp. 3928–3934.
- [4] S. Levine, N. Wagener, and P. Abbeel, "Learning contact-rich manipulation skills with guided policy search," in *Robotics and Automation* (*ICRA*), 2015 IEEE International Conference on. IEEE, 2015, pp. 156–163.
- [5] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proceedings of the IEEE conference on computer* vision and pattern recognition, 2015, pp. 1–9.