Tactile Sensing and Deep Reinforcement Learning for In-Hand Manipulation Tasks

Andrew Melnik[1], Luca Lach[1], Matthias Plappert[2], Timo Korthals[1], Robert Haschke[1], and Helge Ritter[1]

***Abstract*— Deep Reinforcement Learning techniques demonstrate advances in the domain of robotics. One of the limiting factors is the large number of interaction samples usually required for training in simulated and real-world environments. In this work, we demonstrate that tactile information substantially increases sample efficiency for training (by 97% on average), and simultaneously increases the performance in dexterous in-hand manipulation of objects tasks (by 21% on average). To examine the role of tactile-sensor parameters in these improvements, we conducted experiments with varied sensor-measurement accuracy (Boolean vs. float values), and varied spatial resolution of the tactile sensors (92 sensors vs. 16 sensors on the hand). We conclude that ground-truth touch-sensor readings as well as dense tactile resolution do not further improve performance and sample efficiency in the tasks. We make available these touch-sensors extensions as a part of OpenAI-Gym robotics Shadow-Dexterous-Hand environments.**

I. INTRODUCTION

Deep Reinforcement Learning techniques demonstrate advances in the domain of robotics. For example, dexterous in-hand manipulation of objects with an anthropomorphic robotic hand [1, 2]. An agent with a model-free policy was able to learn complex in-hand manipulation tasks using just proprioceptive feedback and visual information about the manipulated object. Continuous haptic feedback can improve grasping acquisition in terms of robustness under uncertainty [3]. We present empirical results in simulation that show that including tactile information in the state improves the sample efficiency and performance for dynamic in-hand manipulation of objects.

Recent works describe approaches to bring the tactile sensing to anthropomorphic hands like the Shadow Dexterous Hand, by providing integrated tactile fingertips [4] as shown in Fig. 1 and constructing a flexible tactile skin [5]. The tactile skin comprises stretchable and flexible, fabric-based tactile sensors capable of capturing typical human interaction forces within the palm and proximal and distal phalanges of the hand. This enables the hand to exploit tactile information, e.g. for contact or slip detection [6]. The distribution of tactile sensors in these works resembles our segmentation of the simulated Shadow Dexterous Hand into 92 tactile-sensitive areas.

II. METHODS

OpenAI Gym [7] contains several simulated robotics environments with the Shadow Dexterous Hand. These environments use the MuJoCo physics engine. The anthropomorphic Shadow Dexterous Hand model, comprising 24 degrees of freedom (20 actuated and 4 coupled), has to manipulate an object (block, egg, or pen) so that it matches a given goal orientation, position, or both position and orientation. For the sake of brevity, further details about training procedure, reward function, goal-aware observation space, and neural network parameters are available in [2], since our main contribution focuses on the extension of the existing Shadow Dexterous Hand model by tactile sensors.

We extended the Shadow Dexterous Hand model with touch sensors available as new environments (Table II) in the OpenAI Gym package [7]. We covered all five fingers

Fig. 1. Shadow Dexterous Hand equipped with fabrics-based tactile sensors in the palm and finger phalanges (indicated green) and fingertip sensors realized by Molded-Interconnect-Devices (indicated yellow) [4, 5]

TABLE I

THE 92 AND 16 TOUCH-SENSOR ENVIRONMENTS.

lower phalanx of the fingers (4x)	7 sensors x 4	1 sensor x 4
middle phalanxes of the fingers (4x)	5 sensors x 4	1 sensor x 4
tip phalanxes of the fingers (4x)	5 sensors x 4	1 sensor x 4
thumb phalanxes (3x)	5 sensors x 3	1 sensor x 3
palm (1x)	9 sensors x 1	1 sensor x 1
All touch sensors	92 sensors	16 sensors

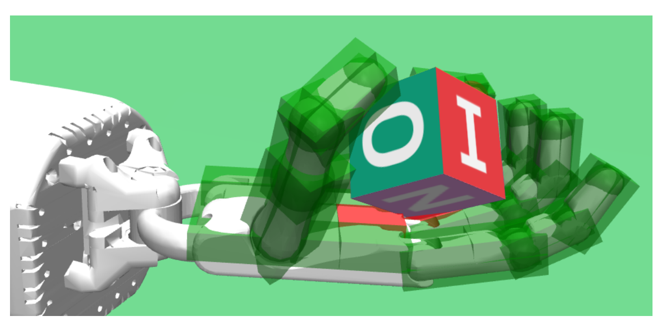

Fig. 2. 92 touch sensors covering the Shadow Dexterous Hand model. This is a technical visualization to represent the essence of our model. Red sites represent activated touch sensors, where a block is pressing against the touch sensitive area. Green sites represent inactive touch sensors. A video demonstration of the extended environments can be found at https://rebrand.ly/TouchSensors

This work was supported by the Cluster of Excellence Cognitive Interaction Technology (CITEC) which is funded by DFG (EXC 277)

[1]CITEC, Bielefeld University, 33619 Bielefeld, Germany. andrew.melnik.papers@gmail.com

[2]OpenAI, San Francisco, USA

TABLE II

OPENAI GYM ROBOTICS ENVIRONMENTS WITH TOUCH SENSORS: -V0 (BOOLEAN), -V1 (FLOAT-VALUE)

HandManipulateBlockRotateZTouchSensors
HandManipulateBlockRotateParallelTouchSensors
HandManipulateBlockRotateXYZTouchSensors
HandManipulateBlockTouchSensors
HandManipulateEggRotateTouchSensors
HandManipulateEggTouchSensors
HandManipulatePenRotateTouchSensors
HandManipulatePenTouchSensors

and the palm of the Shadow Dexterous Hand model with 92 touch sensors (Fig. 2, Table I). For the original OpenAI Gym simulated environments for robotics (Table II) without touch information, the state vector is 68-dimensional [2]. In the environments with 92 touch sensors the state vector is 160-dimensional (68+92). As an additional experiment, we grouped 92 sensors into 16 sub-groups to reduce the tactile sensory resolution ("16 Sensors-v0") (Fig. 2, Table I). In the environments with 16 touch sensors sub-groups the state vector is 84-dimensional (68+16). For a given state of the environment, a trained policy outputs an action vector of 20 float numbers used for position-control (actuation_center + action * actuation_range) of the 20 actuated degrees of freedom.

III. EXPERIMENTAL RESULTS

To provide insights about how different aspects of tactile information (accuracy, tactile resolution) influence learning and performance we conducted three experiments. In the first experiment we added float-value readings from 92 sensors to the state (red curves in Fig. 3). This experiment can be reproduced in the OpenAI-gym-robotics environments ending at "...TouchSensors-v1". In the second experiment we added Boolean-value reading from the same 92 sensors to the state (black curves in Fig. 3). The experiment can be reproduced in the OpenAI-gym-robotics environments ending at "...TouchSensors-v0". In the third experiment we grouped 92 sensors into 16 sub-groups (Table I) to reduce the tactile sensory resolution and added Boolean-value readings from the 16 sub-groups to the state (green curves in Fig. 3). In each experiment, we observe on average 1.21 times better performance (Fig. 3) when tactile information is available. In each experiment, we observe on average 1.97 times faster convergence (Fig. 3) when tactile information is available.

IV. CONCLUSIONS

In this work, we introduce the touch-sensors extensions to OpenAI-Gym [7] robotics Shadow-Dexterous-Hand environments [2] modeled after our touch sensor developments [4, 5]. We find that adding tactile information substantially increases sample efficiency for training (by 97% on average) and performance (by 21% on average) in the environments, when training with deep reinforcement learning techniques. To examine the role of tactile-sensor parameters in these improvements, we conducted experiments (Fig. 3) with varied sensor-measurement accuracy (Boolean vs. float values), and varied spatial resolution of the tactile sensors (92 sensors vs. 16 sensors on the hand). We conclude that accurate sensory readings as well as dense tactile resolution do not substantially improve performance and sample efficiency when training with deep reinforcement learning techniques, in comparison to Boolean sensor readings and sparse sensor localization (one sensor per phalanx). The performance and sample efficiency for training are similar in these case. These experiments provide beneficial knowledge for those looking to build robots with tactile sensors for manipulation.

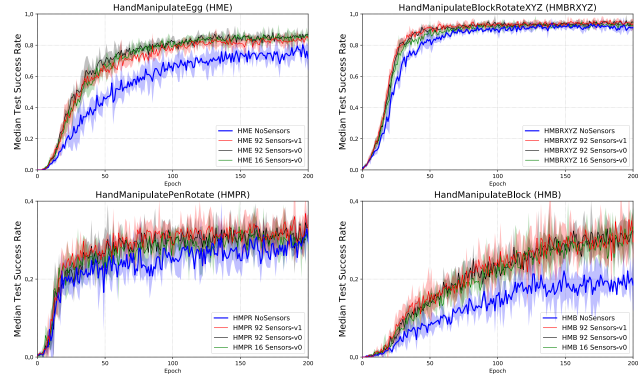

Fig. 3. Curves - median test success rate. Shaded areas - interquartile range [five random seeds]. Blue curves - learning without tactile information (NoSensors), red curves - learning with float-value tactile readings (Sensors-v1) from 92 sensors, black curves - learning with Boolean-value tactile readings (Sensors-v0) from 92 sensors, green curves - learning with Boolean-value tactile readings (Sensors-v0) from 16 sensor sub-groups.

REFERENCES

[1] Marcin Andrychowicz et al. "Learning dexterous in-hand manipulation". In: *arXiv preprint arXiv:1808.00177* (2018).

[2] Matthias Plappert et al. "Multi-goal reinforcement learning: Challenging robotics environments and request for research". In: *arXiv preprint arXiv:1802.09464* (2018).

[3] Timo Korthals et al. "Multisensory Assisted In-hand Manipulation of Objects with a Dexterous Hand". In: *juxi.net* (2019).

[4] Risto Koiva et al. "A highly sensitive 3D-shaped tactile sensor". In: *2013 IEEE/ASME International Conference on Advanced Intelligent Mechatronics*. IEEE. 2013, pp. 1084–1089.

[5] Gereon H Büscher et al. "Flexible and stretchable fabric-based tactile sensor". In: *Robotics and Autonomous Systems* 63 (2015), pp. 244–252.

[6] Martin Meier et al. "Distinguishing sliding from slipping during object pushing". In: *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2016, pp. 5579–5584.

[7] Greg Brockman et al. *OpenAI Gym*. 2016. eprint: arXiv:1606.01540.