TITLE

Generalizing Deep Reinforcement Learning for multi-DoF Robotic Grasping Across Objects

GUIDANCE

• For more information, please contact:	Raffaello Camoriano – raffaello.camoriano@polito.it
• Promoters:	Raffaello Camoriano – Politecnico di Torino Lorenzo Natale – Istituto Italiano di Tecnologia Elisa Maiettini – Istituto Italiano di Tecnologia Federico Ceola – Istituto Italiano di Tecnologia
• Supervisor:	Raffaello Camoriano – Politecnico di Torino Lorenzo Natale – Istituto Italiano di Tecnologia Elisa Maiettini – Istituto Italiano di Tecnologia Federico Ceola – Istituto Italiano di Tecnologia

KEYWORDS

- **Reinforcement Learning**
- **Robot Learning**
- **Grasping and Dexterous Manipulation**
- **Humanoid Robotics**

CONTEXT

Figure 1: The iCub humanoid robot (IIT)

Generating stable object grasps is a longstanding open problem in robotics. The ability of a robot to carry out high-level tasks such as dexterous manipulation [1] hinges on the reliability of the underlying grasping functionality [2]. In turn, grasping difficulty grows significantly with both the number of controlled degrees of freedom (DoFs) of the hand and the geometric complexity and variability of the target objects. In the literature, this problem has been tackled both at the hand control and mechanical design levels, by employing grippers with few degrees of freedom to reduce the complexity of grasp generation or by designing synergical few-DoF soft hands able to passively and robustly adapt to the shape of the object [3].

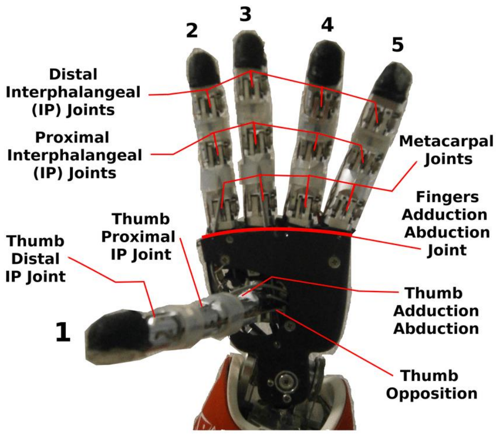

Figure 2: The iCub dexterous hand

However, human-level dexterous manipulation, including in-hand manipulation and grasping of multiple objects, requires hands with many actuated DoFs (a.k.a. multi-DoF hands). This setting results in high-dimensional planning and control problems, further complicated by the large number of hard-to-model contacts between the object and hand parts. Analytical methods employing accurate models to compute optimal contact-rich configurations proved remarkably successful [4]. However, i) high-fidelity models may not be readily available to autonomous robots in unstructured environments, and ii) analytical methods can be prohibitively expensive in terms of computations.

Such limitations motivated the introduction of machine-learning based methods for grasping [5,6]. In particular, model-free deep reinforcement learning (DRL) is a promising framework in this direction [7,8,9], due to its generality and lack of prior modeling requirements. However, an open challenge for DRL is the high dimensionality of the state and action spaces characterizing dexterous multi-DoF hands, resulting in overwhelming computational requirements. This limits the scope of DRL-based grasping policies to few target objects and hampers generality.

REFERENCES

[1] Prattichizzo, Domenico, Maria Pozzi, and Monica Malvezzi. "Dexterous manipulation." Encyclopedia of Robotics. Springer Berlin Heidelberg, Berlin, Heidelberg (2020): 1-8.
[2] Billard, Aude, and Danica Kragic. "Trends and challenges in robot manipulation." Science 364.6446 (2019): eaat8414.
[3] Catalano, Manuel G., et al. "Adaptive synergies for the design and control of the Pisa/IIT SoftHand." The International Journal of Robotics Research 33.5 (2014): 768-782.
[4] Yao, Kunpeng, and Aude Billard. "Exploiting Kinematic Redundancy for Robotic Grasping of Multiple Objects." IEEE Transactions on Robotics (2023). https://arxiv.org/pdf/2303.01868.pdf
[5] Sahbani, Anis, Sahar El-Khoury, and Philippe Bidaud. "An overview of 3D object grasp synthesis algorithms." Robotics and Autonomous Systems 60.3 (2012): 326-336.
[6] Kroemer, Oliver B., et al. "Combining active learning and reactive control for robot grasping." Robotics and Autonomous systems 58.9 (2010): 1105-1116.
[7] Qin, Yuzhe, et al. "DexPoint: Generalizable Point Cloud Reinforcement Learning for Sim-to-Real Dexterous Manipulation." Conference on Robot Learning. PMLR, 2023.
[8] Ghadirzadeh, Ali, et al. "Deep predictive policy training using reinforcement learning." 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 2017.
[9] Andrychowicz, OpenAI: Marcin, et al. "Learning dexterous in-hand manipulation." The International Journal of Robotics Research 39.1 (2020): 3-20.

GOALS

The goal of this thesis is to investigate efficient model-free DRL methods for learning to grasp multiple object categories with multi-DoF robotic hands.

In particular, the candidate will investigate and develop state-of-the-art DRL methods for training multimodal grasping policies capable of generalizing across different object geometries, only given 2D visual input and tactile information. The RL agent will be jointly trained on multiple object categories, exploring the capabilities of various DRL algorithms and multimodal object representations, with a specific focus on improving data efficiency and generalization. The target robotic platform is the iCub humanoid robot with its dexterous multi-DoF hand, resulting in a 20-dimensional action space.

METHODOLOGY

The activity will broadly follow the steps below,:

- Literature review on key model-free DRL algorithms and their application to robotic grasping
- Replication and analysis of state-of-the-art DRL grasping results in simulation, targeting the iCub robot
- Selection, joint training, and empirical evaluation of a DRL agent on multiple object categories
- Investigation of the role of current and alternative multimodal representations on training efficiency, policy generalization, and grasping success rate
- [Optional] Experiments on the real-world iCub robot

PROFILE

Expected load: Full-time, covering a minimum of 6 months, split in:

- Literature and study – 20%
- Implementation – 40%
- Experiments – 40%

Desired skills:

- Good proactivity, teamwork experience, and communication skills are a must;
- Excellent programming and software engineering skills are required;
- Strong experience with Python is required, preferably including state-of-the-art ML, RL, data manipulation, and visualization libraries (e.g., PyTorch, OpenAI Gym, RLLib, stable baselines, mushroom, ...);
- Experience or strong motivation in working with robotics simulators (e.g. Mujoco, Bullet, Gazebo, ...), and potentially with advanced humanoid robots, are welcome;
- Experience with version control (GIT) and experiment management and visualization software (e.g., WandB, TensorBoard, ...) are a plus;
- Foundational knowledge of the Reinforcement learning paradigm is either expected or needs to be gained prior to the start of the thesis;
- Knowledge of robotics fundamentals and/or hands-on experience are a plus

Organization:

- Most activities can be carried out in hybrid mode. The candidate will have access to the VANDAL laboratory premises, with weekly update meetings;
- In-person activities and working sessions at the Center for Robotics and Intelligent

	Systems of Istituto Italiano di Tecnologia in Genoa, where the HSP research group and the iCub robot are physically located, will also be possible. Potential longer stays are optional and conditioned on available resources and interest.
# for Students	1