Doctoral Dissertation
National Doctoral Program in Artificial Intelligence (37^{th}cycle)

Scalable and Generalizable Robot Learning: from Simulated to Real-World Applications

By

Gabriele Tiboni
******

Supervisor(s):
Prof. Tatiana Tommasi, Supervisor
Prof. Jan Peters, Co-Supervisor

Doctoral Examination Committee:
Prof. Jens Kober, Referee
Prof. Joni Pajarinen, Referee
Dr. Mirco Mutti
Dr. Lorenzo Natale
Prof. Alessandro Savino

Politecnico di Torino
2025

Summary

In the past decade, the increased accessibility to high-performance computing power combined with advances in machine learning algorithms has made enormous progress for the development of autonomous systems. In the context of robotics, data-driven methods hold the promise to design intelligent embodied agents that can adapt to previously unseen conditions in real-world environments. Despite the potential of robot learning, however, its practical deployment has for years largely been confined to simulated applications, as learning-based solutions often do not scale to real-world problems due to data scarcity and safety concerns, or are unable to operate in novel conditions. Leveraging physics simulation offers a path to mitigate these challenges by training agents in a safe and efficient manner under a fully controllable environment. Yet, simulated and real-world data distributions exhibit inevitable discrepancies, referred to as the *reality gap*, that heavily hinder the direct transfer of robot behavior to the real world. Similarly, learning complex robotic tasks conditioned on free-form 3D objects also requires models that can transfer their knowledge across domains, *e.g.* for varying shapes and object categories.

This thesis contributes to the challenge of designing intelligent embodied agents, by focusing on algorithms that are applicable to real-world setups and achieve effective generalization across two critical axes: environment dynamics and object geometries.

Part I investigates generalization across dynamics, primarily for sim-to-real transfer of reinforcement learning agents in robot manipulation. We first introduce *DORAEMON*, a zero-shot transfer method building on the maximum entropy principle and curriculum learning to progressively promote stronger generalization capabilities. This approach trains policies in simulation capable of adapting to significant dynamics variations upon real-world deployment without fine-tuning. Subsequently, we present our contributions in the related field of guided domain transfer, with two novel methods that utilize limited, offline, task-agnostic real-world data to automatically tune simulator parameters via statistical inference. Our results show these methods effectively bridge the reality gap with minimal domain knowledge, data constraints and expert supervision, and even for complex soft robot manipulation.

Part II of this thesis tackles generalization across 3D shapes for object-centric motion generation tasks, with in-depth investigation in robotic spray painting. Addressing the lack of available data, we first introduce the *PaintNet* dataset, a large-scale collection of expert demonstrations for painting diverse 3D objects. In this context, we propose *MaskPlanner*, a novel deep learning framework that generates multiple, long-horizon, complex motion paths directly from point cloud representations of the underlying object. MaskPlanner demonstrates

effective generalization to unseen 3D geometries, and produces paths that achieve expert-level paint quality on real objects when deployed on specialized robotic arms, drastically reducing manual programming efforts.

Throughout this thesis, we overall highlight the importance for the community to develop scalable and generalizable robot learning algorithms, with the goal of bridging existing solutions from simulation to real-world applications.