


Cross-Embodiment Movement Primitive Library

Tilman Daab, Sebastian Rietsch, and Tamim Asfour

High Performance Humanoid Technologies Lab (H²T), Karlsruhe Institute of Technology (KIT)

Background. Movement primitives (MPs) [1] are fundamental building blocks that represent basic units of motion, inspired by how humans and animals learn and perform complex actions. They are important in robotics because they enable robots to learn, adapt, and generate a wide variety of behaviors efficiently, often from (human) demonstrations. To enable their reuse, MPs can be stored in a library [2], allowing robots to quickly assemble and modify a set of actions to solve novel tasks. Learning this library incrementally allows to efficiently extend and improve capabilities over time, without relying on storing the original demonstrations permanently [3].



Idea. It would be beneficial to share such libraries across robots, thereby enabling cross-embodiment transfer of skills [4] and making effective use of costly human demonstrations. To achieve this, a MP library can be defined in a robot-agnostic way, e.g., as task-space MPs or as joint-space MPs on a common kinematic structure, and be adapted to different robotic embodiments as needed for execution. The adaptation process, also known as motion retargeting, is a non-trivial problem, as it involves the mapping of movements between disparate kinematic structures with different constraints, degrees of freedom, and physical capabilities, while preserving the functional intent and appearance of the original movement.

Reinforcement Learning (RL) is a promising approach for such retargeting [5], as traditional methods might struggle with human-likeness, respecting self-collisions, or accurately reproducing the end-effector motion. In a typical reinforcement learning application, the training process would produce a policy that allows to efficiently map a canonical representation to an embodiment-specific joint-space representation online. Alternatively, the RL training procedure itself could be used as an optimization procedure, resulting in one (or multiple) translated movement primitive without a reusable policy. Within this proposal, we would like to explore the topic of RL-based motion retargeting, with the aim of learning a MP library across embodiments.

Requirements. The topic is suited for students who enjoy analysing problems deeply and to continue improving the quality and robustness of an approach beyond the first results. Previous experience with reinforcement learning and some robotics background is beneficial. The project will also involve a thorough literature review to identify existing methods and their limitations as a starting point. A publication at a renowned robotics conference like ICRA, Humanoids or IROS is the goal of this project.

[1] Schaal, S., Is Imitation Learning the Route to Humanoid Robots? Trends in Cognitive Sciences, vol. 3, no. 6, pp. 233–242, 1999

[2] Pastor, P., Hoffmann, H., Asfour, T., and Schaal, S., Learning and Generalization of Motor Skills by Learning From Demonstration. IEEE International Conference on Robotics and Automation (ICRA), pp. 763–768, 2009

- [3] Daab, T., Jaquier, N., Dreher, C., Meixner, A., Krebs, F., and Asfour, T., Incremental Learning of Full-Pose Via-Point Movement Primitives on Riemannian Manifolds. IEEE International Conference on Robotics and Automation (ICRA), pp. 2317–2323, 2024
- [4] Liu, J., Li, Z., Yu, M., Dong, Z., Calinon, S., Caldwell, D., and Chen, F., Human–Humanoid Robots' Cross-Embodiment Behavior-Skill Transfer Using Decomposed Adversarial Learning From Demonstration: HOTU, a Human–Humanoid Robots' Skill Transfer Framework. IEEE Robotics & Automation Magazine, vol. 32, no. 1, pp. 68–78, 2025
- [5] Reda, D., Won, J., Ye, Y., van de Panne, M., and Winkler, A., Physics-based motion retargeting from sparse inputs. Proceedings of the ACM on Computer Graphics and Interactive Techniques, vol. 6, no. 3, pp. 1-19, 2023