

## Project Type \_\_\_\_\_

- Master Thesis
- Bachelor Thesis
- Research Project

## Supervisors \_\_\_\_\_

-  Max Nagy
-  nagy@fzi.de
-  Serge Thilges
-  serge.thilges@kit.edu

## Difficulty \_\_\_\_\_

Algorithmic



Math



Application



# The Curse of Dimensionality in Reinforcement Learning

## Motivation

Modern reinforcement learning approaches increasingly operate in high-dimensional spaces. Action chunking methods generate sequences of actions and embodied AI systems—from full-body humanoids to dextrous manipulation—require controlling dozens to hundreds of degrees of freedom. While these high-dimensional action spaces (often paired with equally high-dimensional observation spaces) are known to be challenging, the field currently lacks comprehensive quantitative analysis of how dimensionality itself impacts learning performance. Existing benchmarks conflate task complexity with dimensionality, making it difficult to isolate dimensionality effects and develop principled mitigation strategies. This thesis systematically investigates how observation and action dimensionality affects RL performance, and develops architectural approaches for mitigating dimensionality-induced learning difficulties.



Figure 1: DMC Humanoid, DMC Dog Run and myo hand are examples of high-dim RL tasks. Figure adapted from DIME [1].

## Tasks

### Bachelor Thesis Scope

- Literature Review Survey dimensionality challenges and mitigation strategies in RL.
- Experimental Framework Artificially scale observation/action dimensions using random projection matrices ( $N \times M$  transformations).
- Analysis Quantify relationships between dimensionality and performance (sample efficiency, final performance, stability).

### Master Thesis Scope (extends Bachelor scope)

- Multi-Transition Prediction Implement Actor/Critic architectures operating on multiple environments at once as an alternative way to increase obs/action dimension
- Architectural Mitigations Evaluate architectures for high-dimensional spaces (attention, factored representations, dimension-wise processing).
- Comparative Evaluation Compare mitigation strategies across dimensionality regimes and task complexities.

## Outlook

This work addresses a fundamental open question in reinforcement learning and could provide actionable insights for the design of algorithms targeting high-dimensional control problems. Strong results may lead to publication at a top-tier conference or workshop.

## References

- [1] Onur Celik, Zechu Li, Denis Blessing, Ge Li, Daniel Palenicek, Jan Peters, Georgia Chalvatzaki, and Gerhard Neumann. DIME: Diffusion-based maximum entropy reinforcement learning. In *Forty-second International Conference on Machine Learning*, 2025.