

# The Hidden Linear Structure in Deep Reinforcement Learning

Ezgi Korkmaz, University College London

## Abstract

The utilization of deep neural networks as state-action value function approximators has led to striking progress for reinforcement learning algorithms and applications. Yet the knowledge we have on decision boundary geometry and the loss landscape structure of deep neural policies is still quite limited. In this talk, I will introduce diagnostic tools to analyze and understand the deep reinforcement learning policy manifold.

Orthogonal to the diagnostic perspective, I will describe the adversarial framework proposed in my paper to investigate the decision boundary and deep neural policy manifold similarities across states and across MDPs. Furthermore, I will provide a theoretical analysis that explains that high-sensitivity directions are inevitable in high-dimensions and the foundational reasons behind this phenomenon.

The extensive experiments conducted in the Arcade Learning Environment demonstrate that high-sensitivity directions are shared across states, across MDPs and across algorithms. These results reveal fundamental properties of the deep neural policy manifold structure in high-dimensional state representation MDPs. The theoretical results and the framework introduced shed light onto the structure of the deep neural landscape of artificial intelligence agents that can make decisions in highly complex environments.