Using Artificial Neural Networks for the Automatic Design of Behavioral Arbitrators for Khepera IV Robots

Peter Harmouch

Professor : Alcherio Martinoli
Assistant(s) : Cyrill Baumann

Intelligent beings naturally explore and learn skills, inspiring unsupervised skill discovery in robotics. This field trains robots to acquire skills without explicit rewards, aiming to build a broad library of behaviors to be used in downstream tasks. The objective is to reduce programming time and effort by eliminating the need for unique reward functions for each behavior. Various algorithms have been developed for this purpose and they can be compared through experiments using a Khepera IV robot in a Webots simulation foraging environment.

We focused on two unsupervised deep reinforcement learning algorithms: 'Diversity Is All You Need' (DIAYN), pivotal in skill discovery and previously adapted to our environment in the SP-177 project, and 'Lipschitz-constrained Unsupervised Skill Discovery' (LSD), a more advanced algorithm that we adapted to the same setting. Each algorithm was trained to learn eight skills in a foraging environment featuring a light source and three circular white ground patches.

The skills acquired by DIAYN mainly involve the robot maneuvering in different directions to activate its side light sensors. These skills are rather static, emphasizing basic sensor activations over significant interactions with the environment.

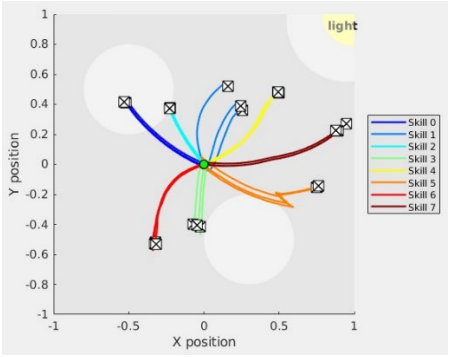

Skills learned by DIAYN

Conversely, the skills developed through LSD demonstrate a more purposeful approach, as each skill leads the robot towards the light source in a unique manner, indicating a higher level of interaction with the environment.

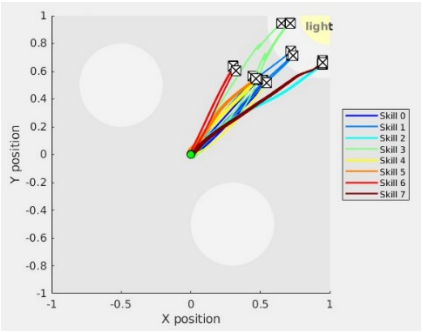

Skills learned by LSD

In a basic task requiring robots to find a randomly positioned light source, LSD's top skill achieved an impressive 99.7% success rate, far surpassing DIAYN's best at 38.4%. Additionally, every skill honed by LSD exhibited unique behaviors in completing the task. Remarkably, one skill skillfully integrated light tracking with ground patch avoidance.

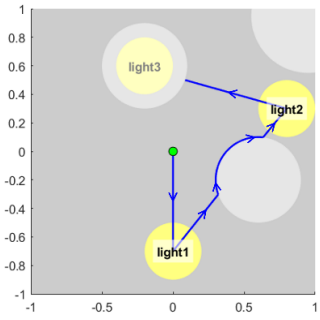

Light following with patch avoidance

It was observed that LSD's effectiveness reduced in complex settings with elements like box obstacles, especially in light-following tasks. To overcome this, implementing Controllability-aware Unsupervised Skill Discovery (CSD), which focuses on harder-to-control state transitions, might be beneficial.