Automatic Design of Behaviors for Khepera IV Robots

Justin Manson

Professor : Alcherio Martinoli
Assistant(s) : Cyrill Baumann

With the increased deployment of mobile robotics in a multitude of real-life scenarios, robots must face increasingly complex and diverse situations. We can develop a hierarchically structured controller to make such robotic platforms possible. Such a system comprises sub-components, each specialized in executing specific tasks, and a behavioral arbitrator that selects the most appropriate sub-component for any given moment. The aforementioned sub-components that solve simple tasks such as obstacle avoidance or light following are today still largely programmed manually. Our work therefore aims to automatically generate a library of basic skills which could then be leveraged using behavioral arbitrators for more complex scenarios.

To this end, we explore two unsupervised deep reinforcement learning algorithms that maximize mutual information between states and skills. Namely, Dynamics-Aware Discovery of Skills (DADS) and in particular Diversity Is All You Need (DIAYN). Both algorithms leverage the mutual information framework to develop a simple reward function that encourages a robot to learn diverse skills.

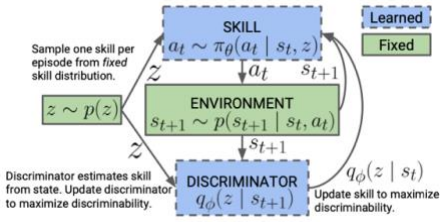

DIAYN algorithm structure

We modify an implementation of DIAYN in order to test the method using webots, a high-fidelity open-source simulation platform. We run a series of experiments using a single Khepera IV robot in multiple variants of a foraging environment. We use various measures of robot state as input for DIAYN and review how these affect the algorithm's ability to converge, as well as the type of skills generated. This leads us to a discussion of the strengths and limitations of the DIAYN algorithm as well as suggestions for axes of further development to improve skill discovery methods.

DIAYN learns skills that lead to diverse states, while DADS, diverse transitions. We believe the latter framework would be more appropriate for goal-oriented tasks such as light following.

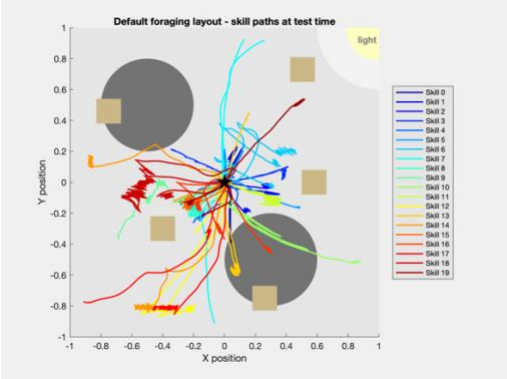

Skills learned in the standard foraging environment

The above example shows paths taken by a library of skills, notably containing a light following and ground foraging behaviors. However, the majority of skills lead to degenerate behaviors that select a very limited set of actions and fail to fully explore the environment. This underscores a general limitation to these skill discovery methods. The mutual information framework ensures that skills are discriminable, however it does not ensure that skills are as diverse as possible and provides no guarantee that they fully explore the environment, or that they are even useful.

We present an alternative framework for unsupervised skill discovery, using reward functions encoded by neural networks, trained iteratively to reward increasingly complex behaviors. This method leads to a DFS-like search and may result in exploration of a particular region in the state space, while neglecting other areas. We therefore suggest that future research initially focus on building a library of simple, diverse skills with DADS and then applying the neural reward function exploration framework to further refine each skill within the library.