Curiosity-Driven Reinforcement Learning with Vision-Language Models

Supervisor(s):
Edward Johns (Department of Computing, e.johns@imperial.ac.uk)

Project description:
When humans are young, we often learn new skills by exploring our environments and discovering new skills all by ourselves. Curiosity-Driven Reinforcement Learning attempts to achieve a similar behaviour with robots. Here, instead of providing a reward function for each new skill, a robot explores its environment autonomously and attempts to perform actions that will result in novel skills being learned, by exploring with "curiosity". For example, if a robot is faced with a new, unfamiliar object that contains a handle, the robot might choose to try pulling on the handle to see what happens, using memories of how "handle-like" parts of objects behaved in the past. Leaving a robot long enough with this kind of exploratory behaviour, might then lead to the robot learning a new, diverse set of skills.

But when a robot learns a new skill, how does it know what this skill is? In this project, you will study how to give a robot the ability to "label" the new skills it is learning, using Vision-Language Models (VLMs) (e.g. GPT-4V). For example, by showing the VLM a few images of the skill being performed, the robot could then generate a language description of the skill. First, you will study and implement the state-of-the-art in Curiosity-Driven Reinforcement Learning, and then, you will study how a VLM can label these new skills so that a robot can learn a language-conditioned policy. In this way, after leaving a robot to explore an environment for a period of time, hopefully the robot can emerge with a set of skills that can be requested by a human, such as "pull open the door", "turn on the tap", and "pick up the apple".

Timeline (tentative):
I am happy for the project to adapt as you progress, according to your own interests and my guidance. But for now, a tentative set of milestones is as follows:
1. Survey the recent literature on curiosity-based RL. (Oct & Nov 2024)
2. Set up a simulation environment (e.g. CoppeliaSim), and use our existing method for imitation learning to generate an initial, basic set of skills for the robot, for simple tabletop tasks, such as picking up objects and placing them elsewhere. (Dec 2024)
3. Introduce curiosity-driven Reinforcement Learning to the simulation, to enable the robot to explore the environment and learn further skills autonomously. (Jan, Feb, Mar 2025)
4. Study how a VLM can be used to label the new skills that have been learned. (Apr & May 2025)
5. (Optional) Evaluate this full framework on more complex tasks in full 3D environments, such as tasks involving opening cupboards, using tools, or performing non-trivial motions such as twisting and swiping. (Jun & Jul 2025)
6. (Optional) Evaluate this full framework in the real world, using one of the robots in our lab. (Jun & Jul 2025)
7. Thesis writing (Aug 2025)

Minimum viable thesis:
The minimum viable thesis would present experimental results up to Step 5 in the proposed timeline. These experiments would show that Curiosity-Driven Reinforcement Learning has been successfully implemented, and that the learned skills can be labelled with a VLM. This would only be in simulation and only with simple tabletop tasks. A more advanced thesis would then include Steps 6 and 7 in the proposed timeline, where experiments show that more complex tasks can be learned, and with real-world experiments.

Required background & skills:

The student should have a good understanding of Reinforcement Learning, good experience with Python programming, and good experience training neural networks, e.g., with PyTorch. Ideally, the student should also be following the latest developments in Robot Learning, particularly new methods that study the intersection of Vision-Language Models and Robotics.

Representative References:

Two example papers on Curiosity-Driven Reinforcement Learning:
https://arxiv.org/pdf/1802.06070.pdf and https://arxiv.org/pdf/2101.04882.pdf
A paper on Curiosity-Driven Reinforcement Learning with a real-world robot:
https://arxiv.org/pdf/2302.06604.pdf
A paper combining Reinforcement Learning with VLMs:
https://arxiv.org/pdf/2307.09668.pdf

Additional notes:

All data for this project will be generated during the project using a robotics simulator.