Deep State Representation Learning and Reinforcement Learning for Robotics

Tutors: David Filliat - david.filliat@ensta-paristech.fr
Natalia Díaz Rodríguez natalia.diaz@ensta-paristech.fr
Inria Flowers Team: https://flowers.inria.fr/ Autonomous Systems and Robotics lab: http://asr.ensta-paristech.fr/ **Unité Informatique et Ingénierie des Systèmes (U2IS)**

Behaviour learning in robotics critically depends on using an appropriate state representation to achieve good performance. For example, to grasp an object, the task is relatively simple if we give as input the 3D position of the object and its exact shape, but much more complex if we simply give an image taken by a camera. However, the definition of the state space requires very good knowledge of the target area and usually a long job done by the designer of the robot. In addition, this state space is often very specific to a problem and can not be reused simply for other tasks. A common example when the robot sensor is a camera is to develop a specific object detection algorithm to achieve the robot's task. In order to obtain robots that are easily adaptable to a changing environment, it would be interesting for them to learn the relevant state space for tasks that are simply displayed by a human operator. The objective is then that the robot automatically detects the data relevant to the task, such as the position of an object or its identity, for example. This area of research called State Representation Learning (SRL)[6] is currently very active through approaches such as Deep Learning that allow to learn representations and features adapted to problems automatically. S-RL Toolbox: Reinforcement Learning (RL) and SRL for Robotics is the developed software library that this internship will use to expand existing methods to improve to more complex scenarios, e.g., multi agent and real robot settings.

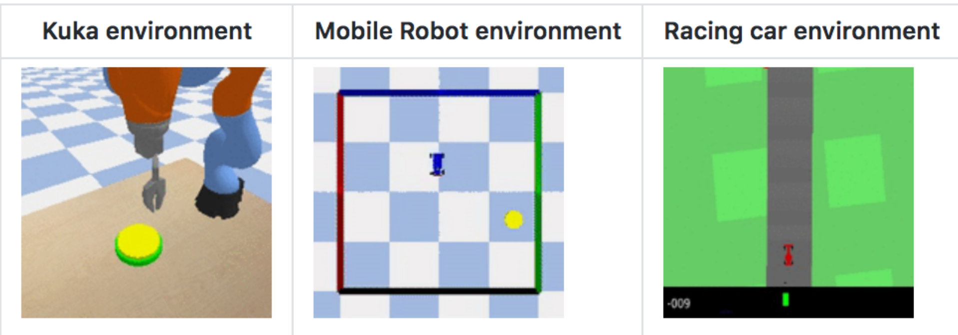

Examples of simulation environments to learn state representations in S-RL Toolbox [1].

The objective of the project will therefore be to implement methods that cope with challenging settings and expand the approaches proposed in [1,3,4] with alternative or complementary strategies such as Generative Adversarial Networks (GANs), or learning from demonstration examples, without having to exactly specify the task. Early work using a Baxter robot [3,4] allowed to validate the approach for simple tasks by using a "Siamese" deep network which makes possible, for example, to implement temporal coherence (among other robotic priors): i.e., *states for two images close in time must be close in the latent space*.
The project will take place at ENSTA ParisTech, in the context of the European project DREAM [2] and will consist of proposing algorithms and evaluating the representation of the states learned (as in [3]) in different environments (simulated / real) through experiments on different state representation learning algorithms. State evaluation can be done in our comparative framework that includes reinforcement learning algorithms or evolving strategies [1], as well as more efficient alternatives to evaluate the quality of the learned representations about robot control tasks. Possible extensions of S-RL Toolbox will be considered in the spirit of continual learning strategies and metrics [5] for avoiding catastrophic forgetting (See ContinualAI.org research community and slack: https://continualai.herokuapp.com/)

References

[1] Raffin et al., 2018. S-RL Toolbox: Environments, Datasets and Evaluation Metrics for State Representation Learning. https://arxiv.org/abs/1809.09369 https://s-rl-toolbox.readthedocs.io https://github.com/araffin/robotics-rl-srl
[2] http://www.robotsthatdream.eu
[3] Decoupling feature extraction from policy learning: assessing benefits of state representation learning in goal based robotics
[4] T. Lesort et al. 2017. Unsupervised state representation learning with robotic priors: a robustness benchmark.
[5] N. Díaz-Rodríguez et al. Don't forget, there is more than forgetting: new metrics for Continual Learning http://arxiv.org/abs/1810.13166
[6] T. Lesort et al. 2018. State Representation Learning for Control: an Overview.https://arxiv.org/abs/1802.04181

Requirements (recommended)

- Image processing, machine learning, deep-learning, reinforcement learning
- Python, PyTorch, TensorFlow, Optionally: PyBullet, ROS, C++