



ellis  
unit

TURIN  
Talk



Politecnico  
di Torino



ARTIFICIAL INTELLIGENCE HUB  
POLITECNICO DI TORINO



<http://vandal.polito.it>

[www.ellis.eu](http://www.ellis.eu)

**Mirco Mutti**

Technion, Israel



## Unsupervised Reinforcement Learning via State Entropy Maximization

In the unsupervised reinforcement learning problem, a policy is first pre-trained from unsupervised interactions, and then fine-tuned towards the optimal policy for several downstream supervised tasks. This recipe has been shown to provide huge benefits w.r.t. learning these tasks from scratch, i.e., starting with a randomly initialized policy, but how to perform the unsupervised pre-training optimally is still an open problem. In this talk, I will focus on the maximum state entropy framework for unsupervised pre-training, in which the agent aims to maximize the entropy over the state visitations induced by its policy. Especially, I will recollect some of the key ingredients that allows to learn maximum state entropy policies in various domains, and I will go through the insights that we got from our work on this specific unsupervised reinforcement learning setting.

**Mirco Mutti** is a postdoctoral researcher working with Aviv Tamar in the Robots Learning Lab at the Technion. Formerly, he was a PhD student at the Università di Bologna and Politecnico di Milano under the supervision of Marcello Restelli. His research focuses on theory and methods for generalization in reinforcement learning, including unsupervised pre-training, partial observability, and meta reinforcement learning.

*April 23rd, 2024*

*starting at 4:00PM*

*Vandal Lab, Covivio - Corso Ferrucci 112  
Turin, Italy*