

CS PhD Seminar Series

July 21st

|

14:30-15:00

|

Room 214

Stochastic Policies for Olfactory Navigation in POMDPs using Natural Actor-Critic

This work investigates whether stochastic policies, optimized via Natural Actor-Critic (NAC) methods, provide superior solutions for olfactory navigation in turbulent environments compared to deterministic Q-learning approaches. Utilizing a compact Finite Clock Memory structure to handle the severe partial observability of chaotic odor plumes, we demonstrate that a stochastic policy parameterization is fundamentally better suited for tracking intermittent chemical signals. Our empirical results show that NAC stochastic policies significantly outperform their deterministic Q-learning counterparts on the identical memory framework, yielding higher target success rates and superior mean trajectory returns. Furthermore, the stochastic formulation dramatically reduces performance variance across trials. By parameterizing action probabilities, the NAC agent naturally executes active, localized casting behaviors that provide robust navigation under intense plume noise where deterministic strategies routinely degrade into cyclic deadlocks.



Speaker: **Abdelmalik Benfadhil**

Abdelmalik Benfadhil is a second-year PhD student at the University of Genoa. He holds both a Bachelor's and a Master's degree in Operations Research from the University of Science and Technology Houari Boumediene (USTHB) in Algiers. During his Master's, he completed a research internship at the International Centre for Theoretical Physics (ICTP). Currently, his research focuses on optimization in reinforcement learning (RL) at the Machine Learning Genoa (MaLGa) research center.