

Learning in input/output systems

Supervised by Bruno Gaujal (Polaris team, Inria and LIG, Grenoble)
Email: bruno.gaujal@inria.fr

Description

Reinforcement learning in partially observable systems focuses on systems where the state is partially hidden by design (such as poker games). Another interesting class which did not attract too much attention up to now is Input/Output systems where the only available observations are inputs and outputs by the internal behavior remains hidden. Unlike general partially observable Markov Decision Processes (POMDPs), these systems can be efficiently learned (see for example [1]), at least for simple control policies, such as admission control. More elaborate policies such as routing are still to be investigated and this will be the object of this internship.

This work could be extended to a three year PhD, for example by investigating a larger set of POMDPs.

Prerequisites

A taste for applied probability and algorithms.

Being knowledgeable in Markov decision processes, linear algebra and stochastic matrices will help.

Information and Contact

This internship duration is 5 to 6 months (M2 internship) and can start anytime in 2024. Contact Bruno Gaujal (bruno.gaujal@inria.fr) for more information.

Bibliography

[1] J. Anselmi, B. Gaujal and J-S Rebuffi. Learning Optimal Admission Control in Partially Observable Queueing Networks, <https://hal.science/hal-04170992/file/questa.pdf>