Bayes-Adaptive Impulse Control of Piecewise-Deterministic Markov Processes

Alice Cleynen[1,2], Benoîte de Saporta[1], Orlane Rossini[1], Régis Sabbadin[3], and Meritxell Vinyals[3]

[1]IMAG, Univ Montpellier, CNRS, Montpellier, France
[2]John Curtin School of Medical Research, Australian National University, Canberra, ACT, Australia
[3]Univ Toulouse, INRAE-MIAT, Toulouse, France

March 20, 2025

1 Abstract

In this paper, we propose a Bayes-adaptive approach for controlling Piecewise-Deterministic Markov Processes (PDMPs) under partial observability and incomplete model knowledge. PDMPs are a form of continuous time (semi-)Markov processes which enable the description of hybrid (discrete-continuous) process dynamics. They allow for the modeling of very general dynamics with a minimal set of interpretable parameters. When some parameters of the controlled PDMP are poorly known, we demonstrate that a Bayes-adaptive approach can provide a *learning while managing* control method. To achieve this, we show how to formulate the problem as a Bayes-Adaptive Partially Observable Markov Decision Process (BAPOMDP), which itself can be modeled as a higher-dimensional Partially Observable Markov Decision Process (POMDP). Subsequently, deep reinforcement learning algorithms can be employed to solve the resulting problem off-line. We illustrate the various steps of our approach with a medical patient follow-up application and utilize the Proximal Policy Optimization (PPO) algorithm to solve the final POMDP model.