
POMDPs in Continuous Time and Discrete Spaces

Bastian Alt* **Matthias Schultheis*,‡** **Heinz Koepl*,‡**

* Department of Electrical Engineering and Information Technology

‡ Centre for Cognitive Science

Technische Universität Darmstadt

{bastian.alt, matthias.schultheis, heinz.koepl}@bcs.tu-darmstadt.de

Abstract

Many processes, such as discrete event systems in engineering or population dynamics in biology, evolve in discrete space and continuous time. We consider the problem of optimal decision making in such discrete state and action space systems under partial observability. This places our work at the intersection of optimal filtering and optimal control. At the current state of research, a mathematical description for simultaneous decision making and filtering in continuous time with finite state and action spaces is still missing. In this paper, we give a mathematical description of a continuous-time partial observable Markov decision process (POMDP). By leveraging optimal filtering theory we derive a Hamilton-Jacobi-Bellman (HJB) type equation that characterizes the optimal solution. Using techniques from deep learning we approximately solve the resulting partial integro-differential equation. We present (i) an approach solving the decision problem offline by learning an approximation of the value function and (ii) an online algorithm which provides a solution in belief space using deep reinforcement learning. We show the applicability on a set of toy examples which pave the way for future methods providing solutions for high dimensional problems.

1 Introduction

Continuous-time models have extensively been studied in machine learning and control. They are especially beneficial to reason about latent variables at time points which are not included in the data. In a broad range of topics such as natural language processing [49], social media dynamics [31] or biology [18] to name just a few, the underlying process naturally evolves continuously in time. In many applications the control of such time-continuous models is of interest. There exist already numerous approaches which tackle the control problem of continuous state space systems, however, for many processes a discrete state space formulation is more suited. This class of systems is discussed in the area of discrete event systems [10]. Decision making in these systems has a long history, yet, if the state is not fully observed acting optimally in such systems is notoriously hard. Many approaches resort to heuristics such as applying a separation principle between inference and control. Unfortunately, this can lead to weak performance as the agent does not incorporate effects of its decisions for future inference.

In the past, this problem was also approached by using a discrete time formulation such as a POMDP model [22]. Nevertheless, it is not always straight-forward to discretize the problem as it requires adding pseudo observations for time points without observations. Additionally, the time discretization can lead to problems when learning optimal controllers in the continuous-time setting [44].

A more principled way to approach this problem is to define the model in continuous time with a proper observation model and to solve the resulting model formulation. Still, it is not clear a priori, how to design such a model and even less how to control it in an optimal way. In this paper, we provide a formulation of this problem by introducing a continuous-time analogue to the POMDP

framework. We additionally show how methods from deep learning can be used to find approximate solutions for control under the continuous-time assumption. Our work can be seen as providing a first step towards scalable control laws for high dimensional problems by making use of further approximation methods. An implementation of our proposed method is publicly available¹.

Related Work

Historically, optimal control in continuous time and space is a classical problem and has been addressed ever since the early works of Pontryagin [35] and Kalman [23]. Continuous-time reinforcement learning formulations have been studied [13, 5, 48] and resulted in works such as the advantage updating and advantage learning algorithms [3, 17] and more recently in function approximation methods [44]. Continuous-time formulations in the case of full observability and discrete spaces are regarded in the context of semi-Markov decision processes (SMDPs) [4, 7], with applications e.g., in E-commerce [53] or as part of the options framework [43, 21].

An important area for the application of time-continuous discrete-state space models is discussed in the queueing theory literature [6]. Here, the state space describes the dynamics of the discrete number of customers in the queue. These models are used for example for internet communication protocols, such as in TCP congestion control. More generally, the topic is also considered within the area of stochastic hybrid systems [11], where mixtures of continuous and discrete state spaces are discussed. Though, the control under partial observability is often only considered under very specific assumptions on the dynamics.

A classical example for simultaneous optimal filtering and control in continuous space is the linear quadratic Gaussian (LQG) controller [41]. In case of partial observability and discrete spaces, the widely used POMDP framework [22] builds a sound foundation for optimal decision making in discrete time and is the basis of many modern methods, e.g., [37, 54, 19]. By applying discretizations, it was also used to solve continuous state space problems as discussed in [12]. Another existing framework which is close to our work is the partial observable semi-Markov decision process (POSMDP) [30] which has applications in fields such as robotics [34] or maintenance [40]. One major drawback of this framework is that observations are only allowed to occur at time points where the latent continuous-time Markov chain (CTMC) jumps. This assumption is actually very limiting, as in many applications observations are received at arbitrary time points or knowledge about jump times is not available.

The development of numerical solution methods for high dimensional partial differential equations (PDEs), such as the HJB equation, is an ongoing topic of research. Popular approaches include techniques based on linearization such as differential dynamic programming [20, 46], stochastic simulation techniques as in the path integral formalism [24, 47] and collocation-based approaches as in [38]. Latter have been extensively discussed due to the recent advances of function approximation by neural networks which have achieved huge empirical success [26]. Approaches solving HJB equations among other PDEs using deep learning can be found in [45, 29, 15, 39].

In the following section we introduce several results from optimal filtering and optimal control, which help to bring our work into perspective.

2 Background

Continuous-Time Markov Chains. First, we discuss the basic notion of the continuous-time Markov chain (CTMC) $\{X(t)\}_{t \geq 0}$, which is a Markov model on discrete state space $\mathcal{X} \subseteq \mathbb{N}$ and in continuous time $t \in \mathbb{R}_{\geq 0}$. Its evolution for a small time step h is characterized by $\mathbb{P}(X(t+h) = x \mid X(t) = x') = \mathbb{1}(x = x') + \Lambda(x', x)h + o(h)$, with $\lim_{h \rightarrow 0} \frac{o(h)}{h} = 0$. Here, $\Lambda : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is the rate function, where $\Lambda(x', x)$ denotes the rate of switching from state x' to state x . We define the exit rate of a state $x \in \mathcal{X}$ as $\Lambda(x) := \sum_{x' \neq x} \Lambda(x, x')$ and $\Lambda(x, x) := -\Lambda(x)$. Due to the memoryless property of the Markov process, the waiting times between two jumps of the chain is exponentially distributed. The definition can also be extended to time dependent rate functions, which correspond to a time inhomogeneous CTMC, for further details see [33].

¹https://git.rwth-aachen.de/bcs/pomdps_continuous_time_discrete_spaces

Optimal Filtering of Latent CTMCs. In the case of partial observability, we do not observe the CTMC $X(t)$ directly, but observe a stochastic process $Y(t)$ depending on $X(t)$. In this setting the goal of optimal filtering [2] is to calculate the filtering distribution, which is also referred to as the belief state at time t ,

$$\pi(x, t) := \mathbb{P}(X(t) = x \mid y_{[0,t]}),$$

where $y_{[0,t]} := \{y(s)\}_{s \in [0,t]}$ is the set of observations in the time interval $[0, t]$. The filtering distribution provides a sufficient statistic to the latent state $X(t)$. We denote the filtering distribution in vector form with components $\{\pi(x, t)\}_{x \in \mathcal{X}}$ by $\pi(t) \in \Delta^{|\mathcal{X}|}$ and use $\Delta^{|\mathcal{X}|}$ for the continuous belief space which is a $|\mathcal{X}|$ dimensional probability simplex.

In general, a large class of filter dynamics follow the description of a jump diffusion process, see [16] for an introduction and the control thereof. The filter dynamics are dependent on an observation model, of which several have been discussed in the literature. A continuous-discrete filter [18] for example uses a model with covariates $\{y_i\}_{i=1, \dots, n}$ observed at discrete time points $t_1 < t_2 < \dots < t_n$, i.e., $y_i := y(t_i)$, according to $y_i \sim p(y_i \mid x)$, where $p(y_i \mid x) := p(y_i \mid X(t_i) = x)$ is a probability density or mass function. In this setting the filtering distribution $\pi(x, t) = \mathbb{P}(X(t) = x \mid y_1, \dots, y_n)$, with $t_n < t$ follows the usual forward equation

$$\frac{d\pi(x, t)}{dt} = \sum_{x'} \Lambda(x', x) \pi(x', t), \quad (1)$$

between observation times t_i and obeys the following reset conditions

$$\pi(x, t_i) = \frac{p(y_i \mid x) \pi(x, t_i^-)}{\sum_{x'} p(y_i \mid x') \pi(x', t_i^-)}, \quad (2)$$

at observation times where $\pi(x, t_i^-)$ denotes the belief just before the observation. The reset conditions at the observation time points represent the posterior distribution of $X(t)$ which combines the likelihood model $p(y \mid x)$ with the previous filtering distribution $\pi(x', t_i^-)$ as prior. The filter equations Eqs. (1) and (2) can be written in the differential form as

$$d\pi(x, t) = \sum_{x'} \Lambda(x', x) \pi(x', t) dt + [\pi(x, t_{N(t)}) - \pi(x, t)] dN(t),$$

which is a special case of a jump diffusion without a diffusion part. The observations enter the filtering equation by the counting process $N(t) = \sum_{i=1}^n \mathbb{1}(t_i \leq t)$ and the jump amplitude $\pi(x, t_{N(t)}) - \pi(x, t)$, which sets the filter to the new posterior distribution $\pi(x, t_{N(t)})$, see Eq. (2).

Other optimal filters can be derived by solving the corresponding Kushner-Stratonovic equation [25, 36]. One instance is the Wonham filter [52, 18] which is discussed in Appendix A.1.

An important observation is that for the case of finite state spaces the filtering distribution is described by a finite dimensional object, opposed to the case of, e.g., uncountable state spaces, where the filtering distribution is described by a probability density, which is infinite dimensional. This is helpful as the filtering distribution can be used as a state in a control setting.

Optimal Control and Reinforcement Learning. Next, we review some results from continuous-time optimal control [41] and reinforcement learning [42], which can be applied to continuous state spaces. Consider a stochastic differential equation (SDE) of the form $d\mathbf{X}(t) = \mathbf{f}(\mathbf{X}(t), u(t))dt + \mathbf{G}(\mathbf{X}(t), u(t))d\mathbf{W}(t)$, where dynamical evolution of the state $\mathbf{X}(t)$ is corrupted by noise following standard Brownian motion $\mathbf{W}(t)$. The control is a function $u : \mathbb{R}_{\geq 0} \rightarrow \mathcal{U}$, where \mathcal{U} is an appropriate action space. In traditional stochastic optimal control the optimal value function, also named cost to go, of a controlled dynamical system is defined by the expected cumulative reward under the optimal policy,

$$V^*(\mathbf{x}) := \max_{u_{[t, \infty)}} \mathbb{E} \left[\int_t^\infty \frac{1}{\tau} e^{-\frac{s-t}{\tau}} R(\mathbf{X}(s), u(s)) ds \mid \mathbf{X}(t) = \mathbf{x} \right], \quad (3)$$

where maximization is carried out over all control trajectories $u_{[t, \infty)} := \{u(s)\}_{s \in [t, \infty)}$. The performance measure or reward function is given by $R : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ and τ denotes the discount factor².

²Note that the discount factor is sometimes defined using certain transformations, e.g., $\frac{1}{\tau} = \rho = -\log \gamma$.

We use a normalization by $1/\tau$ for the value function as it was found to stabilize its learning process when function approximations are used [45]. The optimal value function is given by a PDE, the stochastic HJB equation

$$V^*(\mathbf{x}) = \max_{u \in \mathcal{U}} \left\{ R(\mathbf{x}, u) + \tau \frac{\partial V^*(\mathbf{x})}{\partial \mathbf{x}}^\top \mathbf{f}(\mathbf{x}, u) + \frac{\tau}{2} \text{tr} \left(\frac{\partial^2 V^*(\mathbf{x})}{\partial \mathbf{x}^2} \mathbf{G}(\mathbf{x}, u) \mathbf{G}(\mathbf{x}, u)^\top \right) \right\}.$$

An optimal policy $\mu^* : \mathcal{X} \rightarrow \mathcal{U}$ can be found by solving the PDE and maximizing the right hand side of the HJB equation for every $x \in \mathcal{X}$.

For a deterministic policy $\mu : \mathcal{X} \rightarrow \mathcal{U}$ we define its value as the expected cumulative reward

$$V^\mu(\mathbf{x}) := \mathbb{E} \left[\int_t^\infty \frac{1}{\tau} e^{-\frac{s-t}{\tau}} R(\mathbf{X}(s), U(s)) ds \mid \mathbf{X}(t) = \mathbf{x} \right], \quad (4)$$

where the control $U(s) = \mu(\mathbf{X}(s))$, $s \in [t, \infty)$ follows the policy μ . If the state dynamics are deterministic and the state is differentiable in time, which is the case if there is no diffusion, i.e., $\mathbf{G}(\mathbf{x}, u) = \mathbf{0}$, one can derive a differential equation for the value function [13]. This is achieved by evaluating the value function $V^\mu(\mathbf{x}(t))$ at the function $\mathbf{x}(t)$ and then differentiating both sides of Eq. (4) by time resulting in

$$\tau \frac{d}{dt} V^\mu(\mathbf{x}(t)) = V^\mu(\mathbf{x}(t)) - R(\mathbf{x}(t), \mu(\mathbf{x}(t))). \quad (5)$$

The residuum of Eq. (5) can be identified as a continuous-time analog of the temporal difference (TD)-error

$$\delta(t) := r(t) - V^\mu(\mathbf{x}(t)) + \tau \frac{d}{dt} V^\mu(\mathbf{x}(t)),$$

which an agent should minimize using an observed reward signal $r(t) = R(\mathbf{x}(t), \mu(\mathbf{x}(t)))$ to solve the optimal control problem using reinforcement learning in an online fashion. In this presented case of stochastic optimal control we have assumed that the state is directly usable for the controller. In the partially observable setting, in contrast, this is not the case and the controller can only rely on observations of the state. Hence, the sufficient statistic, the belief state, has to be used for calculating the optimal action. Next, we describe how to model the discrete state space version of the latter case.

3 The Continuous-Time POMDP Model

In this paper we consider a continuous-time equivalent of a POMDP model with time index $t \in \mathbb{R}_{\geq 0}$ defined by the tuple $\langle \mathcal{X}, \mathcal{U}, \mathcal{Y}, \Lambda, \mathbb{P}_{Y|X,u}, R, \tau \rangle$.

We assume a finite state space \mathcal{X} and a finite action space \mathcal{U} . The observation space \mathcal{Y} is either a discrete space or an uncountable space. The latent controlled Markov process follows a CTMC with rate function $\Lambda : \mathcal{X} \times \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$, i.e., $\mathbb{P}(X(t+h) = x \mid X(t) = x', u(t) = u) = \mathbb{1}(x = x') + \Lambda(x', x \mid u)h + o(h)$, with exit rate $\Lambda(x \mid u) = \sum_{x' \neq x} \Lambda(x, x' \mid u)$. Note that the rate function Λ is a function of the control variable. Therefore, the system dynamics are described by a time inhomogeneous CTMC. The underlying process $X(t)$ cannot be directly observed but an observation process $Y(t)$ is available providing information about $X(t)$. The observation model is specified by the conditional probability measure $\mathbb{P}_{Y|X,u}$. The reward function is given by $R : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$. Throughout, we consider an infinite horizon problem with discount τ . We denote the filtering distribution for the latent state $X(t)$ at time t by $\pi(x, t) = \mathbb{P}(X(t) = x \mid y_{[0,t]}, u_{[0,t]})$ with belief state $\boldsymbol{\pi}(t) \in \Delta^{|\mathcal{X}|}$. This filtering distribution is used as state of the partially observable system. Additionally, we define $\mu : \Delta^{|\mathcal{X}|} \rightarrow \mathcal{U}$ as a deterministic policy, which maps a belief state to an action. The performance of a policy μ with control $U(s) = \mu(\boldsymbol{\pi}(s))$, $s \in [0, \infty)$ in the infinite horizon case is given analogously to the standard optimal control case by

$$J^\mu(\boldsymbol{\pi}) = \mathbb{E} \left[\int_0^\infty \frac{1}{\tau} e^{-\frac{s}{\tau}} R(X(s), U(s)) ds \mid \boldsymbol{\pi}(0) = \boldsymbol{\pi} \right],$$

where the expectation is carried out w.r.t. both the latent state and observation processes.

3.1 Simulating the Model

First, we describe a generative process to draw trajectories from a continuous-time POMDP. Consider the stochastic simulation with given policy μ for a trajectory starting at time t with state $X(t) = x'$, and belief $\pi(t)$. The first step is to draw a waiting time ξ of the latent CTMC after which the state switches. This CTMC is time inhomogeneous with exit rate $\Lambda(x' | \mu(\pi(t)))$ since the control modulates the transition function. Therefore, we sample ξ from the time dependent exponential distribution with CDF

$$P(\xi) = 1 - \exp\left(-\int_t^{t+\xi} \Lambda(x' | \mu(\pi(s))) ds\right)$$

for which one needs to solve the filtering trajectory $\{\pi(s)\}_{s \in [t, t+\xi]}$ using a numeric (stochastic) differential equation solver beforehand. There are multiple ways to sample from time dependent exponential distributions. A convenient method to jointly calculate the filtering distribution and the latent state trajectory is provided by the thinning algorithm [27]. For an adaptation for the continuous-time POMDP see Appendix B.1. As second step, after having sampled the waiting time ξ , we can draw the next state $X(t + \xi)$ given ξ from the categorical distribution

$$\mathbb{P}(X(t + \xi) = x | x', \xi) = \begin{cases} \frac{\Lambda(x', x | \mu(\pi(t+\xi)))}{\Lambda(x' | \mu(\pi(t+\xi)))} & \text{if } x' \neq x, \\ 0 & \text{otherwise.} \end{cases}$$

The described process is executed for the initial belief and state at $t = 0$ and repeated until the total simulation time has been reached.

3.2 The HJB Equation in Belief Space

Next, we derive an equation for the value function, which can be solved to obtain an optimal policy. The infinite horizon optimal value function is given by the expected discounted cumulative reward

$$V^*(\pi) = \max_{u_{[t, \infty)}} \mathbb{E} \left[\int_t^\infty \frac{1}{\tau} e^{-\frac{s-t}{\tau}} R(X(s), u(s)) ds \mid \pi(t) = \pi \right]. \quad (6)$$

The value function depends on the belief state π which provides a sufficient statistic for the state. By splitting the integral into two terms from t to $t+h$ and from $t+h$ to ∞ , we have

$$V^*(\pi) = \max_{u_{[t, \infty)}} \mathbb{E} \left[\int_t^{t+h} \frac{1}{\tau} e^{-\frac{s-t}{\tau}} R(X(s), u(s)) ds + \int_{t+h}^\infty \frac{1}{\tau} e^{-\frac{s-t}{\tau}} R(X(s), u(s)) ds \mid \pi(t) = \pi \right]$$

and by identifying the second integral as $e^{-\frac{h}{\tau}} V^*(\pi(t+h))$, we find the stochastic principle of optimality as

$$V^*(\pi) = \max_{u_{[t, t+h)}} \mathbb{E} \left[\int_t^{t+h} \frac{1}{\tau} e^{-\frac{s-t}{\tau}} R(X(s), u(s)) ds + e^{-\frac{h}{\tau}} V^*(\pi(t+h)) \mid \pi(t) = \pi \right]. \quad (7)$$

Here, we consider the class of filtering distributions which follow a jump diffusion process

$$d\pi(t) = \mathbf{f}(\pi(t), u(t))dt + \mathbf{G}(\pi(t), u(t))d\mathbf{W}(t) + \mathbf{h}(\pi(t), u(t))dN(t), \quad (8)$$

where $\mathbf{f} : \Delta^{|\mathcal{X}|} \times \mathcal{U} \rightarrow \mathbb{R}^{|\mathcal{X}|}$ denotes the drift function, $\mathbf{G} : \Delta^{|\mathcal{X}|} \times \mathcal{U} \rightarrow \mathbb{R}^{|\mathcal{X}| \times m}$ the dispersion matrix, with $\mathbf{W}(t) \in \mathbb{R}^m$ being an m dimensional standard Brownian motion and $\mathbf{h} : \Delta^{|\mathcal{X}|} \times \mathcal{U} \rightarrow \mathbb{R}^{|\mathcal{X}|}$ denotes the jump amplitude. We assume a Poisson counting process $N(t) \sim \mathcal{PP}(N(t) | \lambda)$ with rate λ for the observation times $\{t_i\}_{i \in \mathbb{N}}$, which implies that $t_i - t_{i-1} \sim \text{Exp}(t_i - t_{i-1} | \lambda)$. By applying Itô's formula for the jump diffusion processes in Eq. (8) to Eq. (7), dividing both sides by h and taking the limit $h \rightarrow 0$ we find the PDE

$$V^*(\pi) = \max_{u \in \mathcal{U}} \left\{ \mathbb{E}[R(x, u) | \pi] + \tau \frac{\partial V^*(\pi)}{\partial \pi} \mathbf{f}(\pi, u) + \frac{\tau}{2} \text{tr} \left(\frac{\partial^2 V^*(\pi)}{\partial \pi^2} \mathbf{G}(\pi, u) \mathbf{G}(\pi, u)^\top \right) \right. \\ \left. + \tau \lambda (\mathbb{E}[V^*(\pi + \mathbf{h}(\pi, u))] - V^*(\pi)) \right\}, \quad (9)$$

where $\mathbb{E}[R(x, u) | \pi] = \sum_x R(x, u) \pi(x)$. For a detailed derivation see Appendix A.2. We will focus mainly on the case of a controlled continuous-discrete filter

$$d\pi(x, t) = \sum_{x'} \Lambda(x', x | u(t)) \pi(x', t) dt + \left(\frac{p(y_{N(t)} | x, u(t)) \pi(x, t)}{\sum_{x'} p(y_{N(t)} | x', u(t)) \pi(x', t)} - \pi(x, t) \right) dN(t).$$

For this filter, the resulting HJB equation is given by the partial integro-differential equation

$$V^*(\boldsymbol{\pi}) = \max_{u \in \mathcal{U}} \left\{ \sum_x R(x, u) \pi(x) + \tau \sum_{x, x'} \frac{\partial V^*(\boldsymbol{\pi})}{\partial \pi(x)} \Lambda(x', x | u) \pi(x') + \tau \lambda \left(\int \sum_x p(y | x, u) \pi(x) V^*(\boldsymbol{\pi}^\dagger) dy - V^*(\boldsymbol{\pi}) \right) \right\},$$

with $\pi^\dagger(x) = \frac{p(y|x, u)\pi(x)}{\sum_{x'} p(y|x', u)\pi(x')}$. A derivation for the HJB equation for the continuous-discrete filter and for a controlled Wonham filter is provided in Appendix A.3.

To find the optimal policy μ^* corresponding to the optimal value function $V^*(\boldsymbol{\pi})$, it is useful to define the optimal advantage function coinciding with Eq. (9) as

$$A^*(\boldsymbol{\pi}, u) := \mathbb{E}[R(x, u) | \boldsymbol{\pi}] - V^*(\boldsymbol{\pi}) + \tau \frac{\partial V^*(\boldsymbol{\pi})}{\partial \boldsymbol{\pi}} \mathbf{f}(\boldsymbol{\pi}, u) + \frac{\tau}{2} \text{tr} \left(\frac{\partial^2 V^*(\boldsymbol{\pi})}{\partial \boldsymbol{\pi}^2} \mathbf{G}(\boldsymbol{\pi}, u) \mathbf{G}(\boldsymbol{\pi}, u)^\top \right) + \tau \lambda (\mathbb{E}[V^*(\boldsymbol{\pi} + \mathbf{h}(\boldsymbol{\pi}, u))] - V^*(\boldsymbol{\pi})). \quad (10)$$

In order to satisfy Eq. (9), the consistency equation

$$\max_{u \in \mathcal{U}} A^*(\boldsymbol{\pi}, u) = 0 \quad (11)$$

is required. The optimal policy can then be obtained as $\mu^*(\boldsymbol{\pi}) = \arg \max_{u \in \mathcal{U}} A^*(\boldsymbol{\pi}, u)$.

4 Algorithms

For the calculation of an optimal policy we have to find the advantage function. As it depends on the value function, we have to calculate both of them. Since the PDE for the value function in Eq. (9) does not have a closed form solution, we present next two methods to learn the functions approximately.

4.1 Solving the HJB Equation Using a Collocation Method

For solving the HJB equation (9) we first apply a collocation method [38, 45, 29] with a parameterized value function $V_\phi(\boldsymbol{\pi})$, which is modeled by means of a neural network. We define the residual without the maximum operator of the HJB equation under the parameterization as the advantage

$$A_\phi(\boldsymbol{\pi}, u) := \mathbb{E}[R(x, u) | \boldsymbol{\pi}] - V_\phi(\boldsymbol{\pi}) + \tau \frac{\partial V_\phi(\boldsymbol{\pi})}{\partial \boldsymbol{\pi}} \mathbf{f}(\boldsymbol{\pi}, u) + \frac{\tau}{2} \text{tr} \left(\frac{\partial^2 V_\phi(\boldsymbol{\pi})}{\partial \boldsymbol{\pi}^2} \mathbf{G}(\boldsymbol{\pi}, u) \mathbf{G}(\boldsymbol{\pi}, u)^\top \right) + \tau \lambda (\mathbb{E}[V_\phi(\boldsymbol{\pi} + \mathbf{h}(\boldsymbol{\pi}, u))] - V_\phi(\boldsymbol{\pi})).$$

For learning, we sample beliefs $\{\boldsymbol{\pi}^{(i)}\}_{i=1, \dots, N}$, with $\boldsymbol{\pi}^{(i)} \in \Delta^{|\mathcal{X}|}$, from some base distribution $\boldsymbol{\pi}^{(i)} \sim p(\boldsymbol{\pi})$ and estimate the optimal parameters $\hat{\phi} = \arg \min_{\phi} \sum_{i=1}^N \{\max_{u \in \mathcal{U}} A_\phi(\boldsymbol{\pi}^{(i)}, u)\}^2$ by minimizing the squared loss. If needed, we can approximate the expectation $\mathbb{E}[V_\phi(\boldsymbol{\pi}^{(i)} + \mathbf{h}(\boldsymbol{\pi}^{(i)}, u))]$ over the observation space by sampling. An algorithm for this procedure is given in Appendix B.2. For learning it is required to calculate the gradient $\frac{\partial V_\phi(\boldsymbol{\pi})}{\partial \boldsymbol{\pi}}$ and Hessian $\frac{\partial^2 V_\phi(\boldsymbol{\pi})}{\partial \boldsymbol{\pi}^2}$ of the value function w.r.t. the input $\boldsymbol{\pi}$. Generally, this can be achieved by automatic differentiation, but for a fully connected multi-layer feedforward network, the analytic expressions are given in Appendix A.4. The analytic expression makes it possible to calculate the gradient, Hessian and value function in one single forward pass [28].

Given the learned value function $V_{\hat{\phi}}$, we learn an approximation of the advantage function $A_\psi(\boldsymbol{\pi}, u)$ to obtain an approximately-optimal policy. To this end we use a parametrized advantage function $A_\psi(\boldsymbol{\pi}, u)$ and employ a collocation method to solve Eq. (10). To ensure the consistency Eq. (11) during learning, we apply a reparametrization [44, 50] as $A_\psi(\boldsymbol{\pi}, u) = \bar{A}_\psi(\boldsymbol{\pi}, u) - \max_{u' \in \mathcal{U}} \bar{A}_\psi(\boldsymbol{\pi}, u')$. The optimal parameters are found by minimizing the squared loss as

$$\hat{\psi} = \arg \min_{\psi} \sum_{i=1}^N \sum_{u \in \mathcal{U}} \left\{ \bar{A}_\psi(\boldsymbol{\pi}^{(i)}, u) - \max_{u'} \bar{A}_\psi(\boldsymbol{\pi}^{(i)}, u') - A_{\hat{\phi}}(\boldsymbol{\pi}^{(i)}, u) \right\}^2.$$

The corresponding policy can then be easily determined as $\mu(\boldsymbol{\pi}) = \arg \max_{u \in \mathcal{U}} A_{\hat{\psi}}(\boldsymbol{\pi}, u)$ using a single forward pass through the learned advantage function.

4.2 Advantage Updating

The HJB equation (9) can also be solved online using reinforcement learning techniques. We apply the advantage updating algorithm [3] and solve Eq. (10) by employing neural network function approximators for both the value function $V_{\phi}(\boldsymbol{\pi})$ and the advantage function $A_{\psi}(\boldsymbol{\pi}, u)$. The expectations in Eq. (10) are estimated using sample runs of the POMDP. Hence, the residual error can be calculated as

$$E_{\phi, \psi}(t) = A_{\psi}(\boldsymbol{\pi}(t), u(t)) - r(t) + V_{\phi}(\boldsymbol{\pi}(t)) - \tau \frac{\partial V_{\phi}(\boldsymbol{\pi}(t))}{\partial \boldsymbol{\pi}} \mathbf{f}(\boldsymbol{\pi}(t), u(t)) \\ - \frac{\tau}{2} \text{tr} \left(\frac{\partial^2 V_{\phi}(\boldsymbol{\pi}(t))}{\partial \boldsymbol{\pi}^2} \mathbf{G}(\boldsymbol{\pi}(t), u(t)) \mathbf{G}(\boldsymbol{\pi}(t), u(t))^{\top} \right) - \tau \lambda (V_{\phi}(\boldsymbol{\pi}(t^+)) - V_{\phi}(\boldsymbol{\pi}(t))),$$

which can also be seen in terms of a continuous-time TD-error $\delta_{\phi}(t)$ as $E_{\phi, \psi}(t) = A_{\psi}(\boldsymbol{\pi}(t), u(t)) - \delta_{\phi}(t)$. Again we apply the reparametrization $A_{\psi}(\boldsymbol{\pi}, u) = \bar{A}_{\psi}(\boldsymbol{\pi}, u) - \max_{u' \in \mathcal{U}} \bar{A}_{\psi}(\boldsymbol{\pi}, u')$ to satisfy Eq. (11). For estimating the optimal parameters $\hat{\phi}$ and $\hat{\psi}$ we minimize the sum of squared residual errors.

The data for learning is generated by simulating episodes under an exploration policy as in [44]. As exploration policy, we employ a time variant policy $\tilde{\mu}(\boldsymbol{\pi}, t) = \arg \max_{u \in \mathcal{U}} \{A_{\psi}(\boldsymbol{\pi}, u) + \epsilon(u, t)\}$, where $\epsilon(u, t)$ is a stochastic exploration process which we choose as the Ornstein-Uhlenbeck process $d\epsilon(u, t) = -\kappa\epsilon(u, t)dt + \sigma dW(t)$. Generated trajectories are subsampled and saved to a replay buffer [32] which is used to provide data for the training procedure.

5 Experiments

Experimental Tasks. We tested our derived methods on several toy tasks of continuous-time POMDPs with discrete state and action space: An adaption of the popular tiger problem [9], a decentralized multi-agent network transmission problem [8] implementing the slotted aloha protocol, and a grid world problem. All problems are adapted to continuous time and observations at discrete time points. In the following, we provide a brief overview over the considered problems. A more detailed description containing the defined spaces and parameters can be found in Appendix C.

In the tiger problem, the state consists of the position of a tiger (*left/right*) and the agent has to decide between three actions for either improving his belief (*listen*) or exploiting his knowledge to avoid the tiger (*left/right*). While listening the agent can wait for an observation.

In the transmission problem, a set of stations has to adjust their sending rate in order to successfully transmit packages over a single channel as in the slotted Aloha protocol. Each station might have a package to send or not, and the only observation is the past state of the channel, which can be either *idle*, *transmission*, or *collision*. New packages arrive with a fixed rate at the stations. As for each number of available packages – with the exception of no packages available – there is a unique optimal action, a finite number of transmission rates as actions is sufficient.

In the grid world problem, an agent has to navigate through a grid world by choosing the directions in which to move next. The agent transitions with an exponentially distributed amount of time and, while doing so, can slip with some probability so that he instead moves into another direction. The agent receives only noisy information of his position from time to time.

Results. Both the offline collocation method and the online advantage updating method manage to learn reasonable approximations of the value and advantage functions for the considered problems.

For the tiger problem, the learned value and advantage function over the whole belief space are visualized in Fig. 1. The parabolic shape of the value function correctly indicates that higher certainty of the belief about the tiger’s position results in a higher expected discounted cumulative reward as the agent can exploit this knowledge to omit the tiger. Note that the value function learned by the online advantage updating method differs marginally from the one learned by collocation in shape. This is due to the fact that in the advantage updating method only actually visited belief states are used in the learning process to approximate the value function and points in between need to be interpolated.




Figure 1: Learned value and advantage function for the continuous-time tiger problem. The upper plots show the approximated functions learned by the offline collocation method while for lower plots, the online advantage updating method was applied. The orange bars in the lower left plot show the proportions of the beliefs encountered during online learning. The advantage functions on the right can be used to determine the policy of the optimal agent.

The advantage function correctly indicates that for uncertain beliefs it is advantageous to first gain more information by executing the action *listen*. On the other hand, for certain beliefs directly opening the respective door is more useful in terms of reward and opening the opposed door considered as even more disadvantageous in these cases.

Results for the gridworld problem are visualized in Fig. 2 which shows the learned value and advantage function using the online advantage updating method. The figure visualizes the resulting values for certain beliefs, i.e., being at the respective fields with probability one. As expected, the learned value function assigns higher values to fields which are closer to the goal position (3, 2). Actions leading to these states have higher advantage values. For assessing results for uncertain beliefs which are actually encountered when running the system, the figure also contains a sample run which successfully directs the agent to the goal. Results for the collocation method are found to be very similar. A respective plot can be found in Appendix C.2.

For the slotted Aloha transmission problem, a random execution of the policy learned by the offline collocation method is shown in Fig. 3. The upper two plots show the true states, i.e., number of packages and past system state, and the agent’s beliefs which follow the prior dynamics of the system and jump when new observations are made. In the plot at the bottom, the learned advantage function and the resulting policy for the encountered beliefs are visualized. As derived in the appendix, in case of perfect information of the number of packages n , it is optimal to execute action $n - 1$ with exception of $n = 0$ where the action does not matter. When facing uncertainty, however, an optimistic behavior which is also reflected by the learned policy is more reasonable: As for a lower number of packages, the probability of sending a package and therefore collecting higher reward is higher. Thus, in case of uncertainty, one should opt for a lower action than the one executed under perfect information. Results for the online advantage updating method




Figure 2: Learned value and advantage function for certain beliefs in the continuous-time gridworld problem using the advantage updating method. Being at the goal at position (3, 2) results in a reward for the agent. The black fields in row 3 represent a wall that cannot be crossed. Colors of the fields indicate the approximated value function while arrows show the proportions of the advantage functions among different actions. The yellow curvy path indicates a respective random run starting at field (0,0) under the resulting policy.




Figure 3: A sample run for the slotted Aloha transmission problem using a policy learned by the collocation method. The upper plot shows the actual number of packages available in the system while the middle one shows the past system state which can be either *idle* (i), *transmission* (t), or *collision* (c). The background of both plots indicate the marginal belief of number of packages and past system state, respectively. The past system states that are observed at discrete time points are indicated by a cross. The lower plot shows the policy of the agent resulting from the learned advantage function while the per timestep normalized advantage function is indicated by the background.

look qualitatively equal reflecting the same reasonable behavior. A respective plot is omitted here but can be found in Appendix C.2.

6 Conclusion

In this work we introduced a new model for decision making in continuous time under partial observability. We presented (i) a collocation-based offline method and (ii) an advantage updating online algorithm, which both find approximate solutions by the use of deep learning techniques. For evaluation we discussed qualitatively the found solutions for a set of toy problems that were adapted from literature. The solutions have shown to represent reasonable optimal decisions for the given problems.

In the future we are interested in exploring ways for making the proposed model applicable to more realistic large-scale problems. First, throughout this work we made the assumption of a known model. In many applications, however, this might not be the case and investigating how to realize dual control methods [14, 1] might be a fruitful direction. New scalable techniques for estimating parameters of latent CTMCs which could be used, are discussed in [51] but also learning the filtering distribution directly from data might be an option if the model is not available [19]. An issue we faced, was that for high dimensional problems the algorithms seemed to slowly converge to the optimal solution as the belief space grows linearly in the number of dimensions w.r.t. the number of states of the latent process. The introduction of variational and sampling methods seems to be promising to project the filtering distribution to a lower dimensional space and make solution of high-dimensional problems feasible. This could also enable the extension to continuous state spaces, where the filtering distribution is in general infinite dimensional and has to be projected to a finite dimensional representation, e.g., as it is done in assumed density filtering [36]. This will enable the use for interesting applications for example in queuing networks or more generally, in stochastic hybrid systems.

Acknowledgments

This work has been funded by the German Research Foundation (DFG) as part of the project B4 within the Collaborative Research Center (CRC) 1053 – MAKI, and by the European Research Council (ERC) within the CONSYN project, grant agreement number 773196.

Broader Impact

Not applicable to this manuscript.

References

- [1] B. Alt, A. Šošić, and H. Koepl. Correlation priors for reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 14155–14165, 2019.
- [2] A. Bain and D. Crisan. *Fundamentals of stochastic filtering*, volume 60. Springer Science & Business Media, 2008.
- [3] L. C. Baird. Reinforcement learning in continuous time: Advantage updating. In *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94)*, volume 4, pages 2448–2453. IEEE, 1994.
- [4] D. P. Bertsekas. *Dynamic programming and optimal control*, volume 2. Athena scientific Belmont, MA, 1995.
- [5] S. Bhasin, R. Kamalapurkar, M. Johnson, K. G. Vamvoudakis, F. L. Lewis, and W. E. Dixon. A novel actor-critic-identifier architecture for approximate optimal control of uncertain nonlinear systems. *Automatica*, 49(1):82–92, 2013.
- [6] G. Bolch, S. Greiner, H. De Meer, and K. S. Trivedi. *Queueing networks and Markov chains: modeling and performance evaluation with computer science applications*. John Wiley & Sons, 2006.
- [7] S. J. Bradtke and M. O. Duff. Reinforcement learning methods for continuous-time Markov decision problems. In *Advances in neural information processing systems*, pages 393–400, 1995.
- [8] A. R. Cassandra. *Exact and Approximate Algorithms for Partially Observable Markov Decision Processes*. PhD thesis, Department of Computer Science, Brown University, Providence, RI, 1998.
- [9] A. R. Cassandra, L. P. Kaelbling, and M. L. Littman. Acting optimally in partially observable stochastic domains. In *AAAI*, volume 94, pages 1023–1028, 1994.
- [10] C. G. Cassandras and S. Lafortune. *Introduction to discrete event systems*. Springer Science & Business Media, 2009.
- [11] C. G. Cassandras and J. Lygeros. *Stochastic hybrid systems*. CRC Press, 2018.
- [12] P. Chaudhuri, S. Karaman, D. Hsu, and E. Frazzoli. Sampling-based algorithms for continuous-time POMDPs. In *2013 American Control Conference*, pages 4604–4610. IEEE, 2013.
- [13] K. Doya. Reinforcement learning in continuous time and space. *Neural computation*, 12(1):219–245, 2000.
- [14] A. Feldbaum. Dual control theory. i. *Avtomatika i Telemekhanika*, 21(9):1240–1249, 1960.
- [15] J. Han, A. Jentzen, and E. Weinan. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- [16] F. B. Hanson. *Applied stochastic processes and control for jump-diffusions: modeling, analysis and computation*. SIAM, 2007.
- [17] M. E. Harmon and L. C. Baird III. Multi-player residual advantage learning with general function approximation. *Wright Laboratory, WL/AACF, Wright-Patterson Air Force Base, OH*, pages 45433–7308, 1996.
- [18] L. Huang, L. Pauleve, C. Zechner, M. Unger, A. S. Hansen, and H. Koepl. Reconstructing dynamic molecular states from single-cell time series. *Journal of The Royal Society Interface*, 13(122):20160533, 2016.
- [19] M. Igl, L. Zintgraf, T. A. Le, F. Wood, and S. Whiteson. Deep variational reinforcement learning for POMDPs. In *International Conference on Machine Learning*, pages 2117–2126, 2018.
- [20] D. H. Jacobson. New second-order and first-order algorithms for determining optimal control: A differential dynamic programming approach. *Journal of Optimization Theory and Applications*, 2(6):411–440, 1968.

- [21] A. Jain and D. Precup. Eligibility traces for options. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1008–1016. International Foundation for Autonomous Agents and Multiagent Systems, 2018.
- [22] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.
- [23] R. Kalman. The theory of optimal control and the calculus of variations. In *Mathematical optimization techniques*, pages 309–331. University of California Press Los Angeles, CA, 1963.
- [24] H. J. Kappen. Path integrals and symmetry breaking for optimal control theory. *Journal of statistical mechanics: theory and experiment*, 2005(11):P11011, 2005.
- [25] H. J. Kushner. On the differential equations satisfied by conditional probability densities of Markov processes, with applications. *Journal of the Society for Industrial and Applied Mathematics, Series A: Control*, 2(1):106–119, 1964.
- [26] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.
- [27] P. W. Lewis and G. S. Shedler. Simulation of nonhomogeneous Poisson processes by thinning. *Naval research logistics quarterly*, 26(3):403–413, 1979.
- [28] M. Lutter, C. Ritter, and J. Peters. Deep Lagrangian networks: Using physics as model prior for deep learning. In *International Conference on Learning Representations*, 2019.
- [29] M. Lutter, B. Belousov, K. Listmann, D. Clever, and J. Peters. HJB optimal feedback control with deep differential value functions and action constraints. In *Proceedings of the Conference on Robot Learning*, pages 640–650. PMLR, 30 Oct–01 Nov 2020.
- [30] S. Mahadevan. Partially observable semi-Markov decision processes: Theory and applications in engineering and cognitive science. In *AAAI Fall Symposium on Planning with Partially Observable Markov Decision Processes, Orlando, FL, USA*, pages 113–120, 1998.
- [31] H. Mei and J. M. Eisner. The neural Hawkes process: A neurally self-modulating multivariate point process. In *Advances in Neural Information Processing Systems*, pages 6754–6764, 2017.
- [32] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [33] J. R. Norris. *Markov chains*. Number 2 in Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge university press, 1998.
- [34] S. Omidshafiei, A.-A. Agha-Mohammadi, C. Amato, S.-Y. Liu, J. P. How, and J. Vian. Graph-based cross entropy method for solving multi-robot decentralized POMDPs. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5395–5402. IEEE, 2016.
- [35] L. S. Pontryagin. *The mathematical theory of optimal processes*. John Wiley, 1962.
- [36] S. Särkkä and A. Solin. *Applied stochastic differential equations*, volume 10. Cambridge University Press, 2019.
- [37] D. Silver and J. Veness. Monte-carlo planning in large POMDPs. In *Advances in neural information processing systems*, pages 2164–2172, 2010.
- [38] A. Simpkins and E. Todorov. Practical numerical methods for stochastic optimal control of biological systems in continuous time and space. In *2009 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning*, pages 212–218. IEEE, 2009.
- [39] J. Sirignano and K. Spiliopoulos. DGM: A deep learning algorithm for solving partial differential equations. *Journal of Computational Physics*, 375:1339–1364, 2018.
- [40] R. Srinivasan and A. K. Parlikad. Semi-Markov decision process with partial information for maintenance decisions. *IEEE Transactions on Reliability*, 63(4):891–898, 2014.
- [41] R. F. Stengel. *Optimal control and estimation*. Courier Corporation, 1994.
- [42] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT Press, 2018.

- [43] R. S. Sutton, D. Precup, and S. Singh. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2):181–211, 1999.
- [44] C. Tallec, L. Blier, and Y. Ollivier. Making deep q-learning methods robust to time discretization. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 6096–6104, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- [45] Y. Tassa and T. Erez. Least squares solutions of the HJB equation with neural network value-function approximators. *IEEE transactions on neural networks*, 18(4):1031–1041, 2007.
- [46] Y. Tassa, T. Erez, and W. D. Smart. Receding horizon differential dynamic programming. In *Advances in neural information processing systems*, pages 1465–1472, 2008.
- [47] E. A. Theodorou and E. Todorov. Stochastic optimal control for nonlinear Markov jump diffusion processes. In *2012 American Control Conference (ACC)*, pages 1633–1639. IEEE, 2012.
- [48] K. G. Vamvoudakis and F. L. Lewis. Online actor–critic algorithm to solve the continuous-time infinite horizon optimal control problem. *Automatica*, 46(5):878–888, 2010.
- [49] C. Wang, D. Blei, and D. Heckerman. Continuous time dynamic topic models. In *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence*, pages 579–586, 2008.
- [50] Z. Wang, T. Schaul, M. Hessel, H. Van Hasselt, M. Lanctot, and N. De Freitas. Dueling network architectures for deep reinforcement learning. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning-Volume 48*, pages 1995–2003, 2016.
- [51] C. Wildner and H. Koepl. Moment-based variational inference for Markov jump processes. In *International Conference on Machine Learning*, pages 6766–6775, 2019.
- [52] W. M. Wonham. Some applications of stochastic differential equations to optimal nonlinear filtering. *Journal of the Society for Industrial and Applied Mathematics, Series A: Control*, 2(3):347–369, 1964.
- [53] H. Xie, Y. Li, and J. C. Lui. Optimizing discount and reputation trade-offs in e-commerce systems: Characterization and online learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7992–7999, 2019.
- [54] N. Ye, A. Somani, D. Hsu, and W. S. Lee. Despot: Online POMDP planning with regularization. *Journal of Artificial Intelligence Research*, 58:231–266, 2017.

POMDPs in Continuous Time and Discrete Spaces

— Supplementary Material —

Bastian Alt* **Matthias Schultheis*,[‡]** **Heinz Koepl*,[‡]**

* Department of Electrical Engineering and Information Technology

[‡] Centre for Cognitive Science

Technische Universität Darmstadt

{bastian.alt, matthias.schultheis, heinz.koepl}@bcs.tu-darmstadt.de

A Models and Derivations

A.1 The Wonham filter

Another well known observation model is

$$d\mathbf{Y}(t) = \mathbf{g}(\mathbf{X}(t), t)dt + \mathbf{B}(t)d\mathbf{W}(t).$$

Here, the l -dimensional observation process $\mathbf{Y}(t) \in \mathbb{R}^l$ is described by an SDE with a drift function $\mathbf{g} : \mathcal{X} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^l$ depending on the latent state of the CTMC. Noise is added by the dispersion matrix $\mathbf{B}(t) \in \mathbb{R}^{l \times m}$ and the Brownian motion $\mathbf{W}(t) \in \mathbb{R}^m$. The filtering distribution, here, can be described by the Wonham filter [52, 18]

$$d\pi(x, t) = \sum_{x'} \Lambda(x', x)\pi(x', t)dt + \pi(x, t)(\mathbf{g}(x, t) - \bar{\mathbf{g}}(t))^\top (\mathbf{B}(t)\mathbf{B}(t)^\top)^{-1} (d\mathbf{Y}(t) - \bar{\mathbf{g}}(t)dt),$$

where $\bar{\mathbf{g}}(t) = \sum_{x'} \mathbf{g}(x', t)\pi(x', t)$.

A.2 The HJB Equation for a Filter Given by a Jump Diffusion Process

The principle of optimality reads

$$V^*(\boldsymbol{\pi}) = \max_{u_{[t, t+h]}} \mathbb{E} \left[\int_t^{t+h} \frac{1}{\tau} e^{-\frac{s-t}{\tau}} R(X(s), u(s))ds + e^{-\frac{h}{\tau}} V^*(\boldsymbol{\pi}(t+h)) \mid \boldsymbol{\pi}(t) = \boldsymbol{\pi} \right]. \quad (12)$$

Assuming a jump diffusion dynamic as

$$d\boldsymbol{\pi}(t) = \mathbf{f}(\boldsymbol{\pi}(t), u(t))dt + \mathbf{G}(\boldsymbol{\pi}(t), u(t))d\mathbf{W}(t) + \mathbf{h}(\boldsymbol{\pi}(t), u(t))dN(t),$$

we apply Itô's formula for jump diffusion processes to the value function and find

$$\begin{aligned} V^*(\boldsymbol{\pi}(t+h)) &= V^*(\boldsymbol{\pi}(t)) + \int_t^{t+h} \frac{\partial V^*(\boldsymbol{\pi}(s))}{\partial \boldsymbol{\pi}} \mathbf{f}(\boldsymbol{\pi}(s), u(s))ds \\ &\quad + \int_t^{t+h} \frac{\partial V^*(\boldsymbol{\pi}(s))}{\partial \boldsymbol{\pi}} \mathbf{G}(\boldsymbol{\pi}(s), u(s))d\mathbf{W}(s) \\ &\quad + \int_t^{t+h} \frac{1}{2} \operatorname{tr} \left(\frac{\partial^2 V^*(\boldsymbol{\pi}(s))}{\partial \boldsymbol{\pi}^2} \mathbf{G}(\boldsymbol{\pi}(s), u(s))\mathbf{G}(\boldsymbol{\pi}(s), u(s))^\top \right) ds \\ &\quad + \int_t^{t+h} V^*(\boldsymbol{\pi}(s) + \mathbf{h}(\boldsymbol{\pi}(s), u(s)) - V^*(\boldsymbol{\pi}(s))dN(s). \end{aligned} \quad (13)$$

By inserting Eq. (13) into Eq. (12) we find

$$\begin{aligned}
V^*(\boldsymbol{\pi}) &= \max_{u_{[t, t+h]}} \mathbb{E} \left[\int_t^{t+h} \frac{1}{\tau} e^{-\frac{s-t}{\tau}} R(X(s), u(s)) ds + e^{-\frac{h}{\tau}} (V^*(\boldsymbol{\pi}(t))) \right. \\
&\quad + \int_t^{t+h} \frac{\partial V^*(\boldsymbol{\pi}(s))}{\partial \boldsymbol{\pi}} \mathbf{f}(\boldsymbol{\pi}(s), u(s)) ds + \int_t^{t+h} \frac{\partial V^*(\boldsymbol{\pi}(s))}{\partial \boldsymbol{\pi}} \mathbf{G}(\boldsymbol{\pi}(s), u(s)) d\mathbf{W}(s) \\
&\quad + \int_t^{t+h} \frac{1}{2} \text{tr} \left(\frac{\partial^2 V^*(\boldsymbol{\pi}(s))}{\partial \boldsymbol{\pi}^2} \mathbf{G}(\boldsymbol{\pi}(s), u(s)) \mathbf{G}(\boldsymbol{\pi}(s), u(s))^\top \right) ds \\
&\quad \left. + \int_t^{t+h} V^*(\boldsymbol{\pi}(s) + \mathbf{h}(\boldsymbol{\pi}(s)), u(s)) - V^*(\boldsymbol{\pi}(s)) dN(s) \right] \Big| \boldsymbol{\pi}(t) = \boldsymbol{\pi}.
\end{aligned}$$

Collecting terms in $V^*(\boldsymbol{\pi})$ and dividing both sides by h we get

$$\begin{aligned}
V^*(\boldsymbol{\pi}) \frac{1 - e^{-\frac{h}{\tau}}}{h} &= \max_{u_{[t, t+h]}} \mathbb{E} \left[\frac{1}{h} \int_t^{t+h} \frac{1}{\tau} e^{-\frac{s-t}{\tau}} R(X(s), u(s)) ds + \frac{e^{-\frac{h}{\tau}}}{h} \left(\int_t^{t+h} \frac{\partial V^*(\boldsymbol{\pi}(s))}{\partial \boldsymbol{\pi}} \mathbf{f}(\boldsymbol{\pi}(s), u(s)) ds + \int_t^{t+h} \frac{\partial V^*(\boldsymbol{\pi}(s))}{\partial \boldsymbol{\pi}} \mathbf{G}(\boldsymbol{\pi}(s), u(s)) d\mathbf{W}(s) \right. \right. \\
&\quad + \int_t^{t+h} \frac{1}{2} \text{tr} \left(\frac{\partial^2 V^*(\boldsymbol{\pi}(s))}{\partial \boldsymbol{\pi}^2} \mathbf{G}(\boldsymbol{\pi}(s), u(s)) \mathbf{G}(\boldsymbol{\pi}(s), u(s))^\top \right) ds \\
&\quad \left. \left. + \int_t^{t+h} V^*(\boldsymbol{\pi}(s) + \mathbf{h}(\boldsymbol{\pi}(s)), u(s)) - V^*(\boldsymbol{\pi}(s)) dN(s) \right) \Big| \boldsymbol{\pi}(t) = \boldsymbol{\pi} \right].
\end{aligned}$$

Taking $\lim_{h \rightarrow 0}$ and calculating the expectation w.r.t. $\mathbf{W}(t)$ and $N(t)$ we find

$$\begin{aligned}
\frac{1}{\tau} V^*(\boldsymbol{\pi}) &= \max_u \frac{1}{\tau} \mathbb{E} [R(X(t), u) | \boldsymbol{\pi}(t) = \boldsymbol{\pi}] + \frac{\partial V^*(\boldsymbol{\pi})}{\partial \boldsymbol{\pi}} \mathbf{f}(\boldsymbol{\pi}, u) \\
&\quad + \frac{1}{2} \text{tr} \left(\frac{\partial^2 V^*(\boldsymbol{\pi})}{\partial \boldsymbol{\pi}^2} \mathbf{G}(\boldsymbol{\pi}, u) \mathbf{G}(\boldsymbol{\pi}, u)^\top \right) + \lambda (\mathbb{E} [V^*(\boldsymbol{\pi} + \mathbf{h}(\boldsymbol{\pi}))] - V^*(\boldsymbol{\pi})),
\end{aligned}$$

for further details, see [16].

A.3 The HJB Equation for the Presented Filters

A.3.1 The HJB Equation for the Continuous Discrete Filter

The HJB equation in belief space is

$$\begin{aligned}
V^*(\boldsymbol{\pi}) &= \max_{u \in \mathcal{U}} \left\{ \mathbb{E} [R(x, u) | \boldsymbol{\pi}] + \tau \frac{\partial V^*(\boldsymbol{\pi})}{\partial \boldsymbol{\pi}} \mathbf{f}(\boldsymbol{\pi}, u) + \frac{\tau}{2} \text{tr} \left(\frac{\partial^2 V^*(\boldsymbol{\pi})}{\partial \boldsymbol{\pi}^2} \mathbf{G}(\boldsymbol{\pi}, u) \mathbf{G}(\boldsymbol{\pi}, u)^\top \right) \right. \\
&\quad \left. + \tau \lambda (\mathbb{E} [V^*(\boldsymbol{\pi} + \mathbf{h}(\boldsymbol{\pi}, u))] - V^*(\boldsymbol{\pi})) \right\}.
\end{aligned}$$

The filter dynamics for the controlled continuous-discrete filter are given by

$$d\pi(x, t) = \sum_{x'} \Lambda(x', x | u(t)) \pi(x', t) dt + \left(\frac{p(y_{N(t)} | x, u(t)) \pi(x, t)}{\sum_{x'} p(y_{N(t)} | x', u(t)) \pi(x', t)} - \pi(x, t) \right) dN(t).$$

Hence, the components $\{f(\boldsymbol{\pi}, u, x)\}_{x \in \mathcal{X}}$ of the drift function $f(\boldsymbol{\pi}, u)$ are given by

$$f(\boldsymbol{\pi}, u, x) = \sum_{x'} \Lambda(x', x | u) \pi(x')$$

and the diffusion term is zero, i.e., $\mathbf{G}(\boldsymbol{\pi}, u) = \mathbf{0}$. The components $\{h(\boldsymbol{\pi}, u, x)\}_{x \in \mathcal{X}}$ of the jump amplitude $\mathbf{h}(\boldsymbol{\pi}, u)$ are

$$h(\boldsymbol{\pi}, u, x) = \frac{p(y | x, u) \pi(x)}{\sum_{x'} p(y | x', u) \pi(x')} - \pi(x),$$

with $y \sim p(y)$ and $p(y) = \sum_x p(y | x, u)\pi(x)$. Thus, the expectation yields

$$\mathbb{E}[V^*(\boldsymbol{\pi} + \mathbf{h}(\boldsymbol{\pi}, u))] = \int \sum_x p(y | x, u)\pi(x)V^*(\boldsymbol{\pi}^+)dy,$$

with $\boldsymbol{\pi}^+(x) = \frac{p(y|x, u)\pi(x)}{\sum_{x'} p(y|x', u)\pi(x')}$. Finally, we have the HJB equation

$$V^*(\boldsymbol{\pi}) = \max_{u \in \mathcal{U}} \left\{ \sum_x R(x, u)\pi(x) + \tau \sum_{x, x'} \frac{\partial V^*(\boldsymbol{\pi})}{\partial \pi(x)} \Lambda(x', x | u)\pi(x') + \tau \lambda \left(\int \sum_x p(y | x, u) \pi(x) V^*(\boldsymbol{\pi}^+) dy - V^*(\boldsymbol{\pi}) \right) \right\}.$$

A.3.2 The HJB Equation for the Wonham Filter

We consider the HJB equation in belief space

$$V^*(\boldsymbol{\pi}) = \max_{u \in \mathcal{U}} \left\{ \mathbb{E}[R(x, u) | \boldsymbol{\pi}] + \tau \frac{\partial V^*(\boldsymbol{\pi})}{\partial \boldsymbol{\pi}} \mathbf{f}(\boldsymbol{\pi}, u) + \frac{\tau}{2} \text{tr} \left(\frac{\partial^2 V^*(\boldsymbol{\pi})}{\partial \boldsymbol{\pi}^2} \mathbf{G}(\boldsymbol{\pi}, u) \mathbf{G}(\boldsymbol{\pi}, u)^\top \right) + \tau \lambda (\mathbb{E}[V^*(\boldsymbol{\pi} + \mathbf{h}(\boldsymbol{\pi}, u))] - V^*(\boldsymbol{\pi})) \right\}$$

and the controlled Wonham filter

$$d\boldsymbol{\pi}(x, t) = \sum_{x'} \Lambda(x', x | u(t))\pi(x', t)dt + \pi(x, t)(\mathbf{g}(x, u(t)) - \bar{\mathbf{g}}(u(t), t)(\mathbf{B}\mathbf{B}^\top)^{-1}(d\mathbf{Y}(t) - \bar{\mathbf{g}}(u(t), t)dt),$$

with $\bar{\mathbf{g}}(u(t), t) = \sum_{x'} \mathbf{g}(x', u(t))\pi(x', t)$. Therefore, we have the components $\{f(\boldsymbol{\pi}, u, x)\}_{x \in \mathcal{X}}$ of the drift function $\mathbf{f}(\boldsymbol{\pi}, u)$ as

$$f(\boldsymbol{\pi}, u, x) = \sum_{x'} \Lambda(x', x | u)\pi(x') + \pi(x)(\mathbf{g}(x, u) - \bar{\mathbf{g}}(u)^\top (\mathbf{B}\mathbf{B}^\top)^{-1}(\mathbf{g}(x, u) - \bar{\mathbf{g}}(u))).$$

The row-wise components $\{\mathbf{G}(\boldsymbol{\pi}, u, x)\}_{x \in \mathcal{X}}$ of the dispersion matrix $\mathbf{G}(\boldsymbol{\pi}, u)$ read

$$\mathbf{G}(\boldsymbol{\pi}, u, x) = \pi(x)\mathbf{B}^\top (\mathbf{B}\mathbf{B}^\top)^{-1}(\mathbf{g}(x, u) - \bar{\mathbf{g}}(u)),$$

with $\bar{\mathbf{g}}(u) = \sum_{x'} \mathbf{g}(x', u)\pi(x')$. The jump amplitude is zero, i.e., $\mathbf{h}(\boldsymbol{\pi}, u) = \mathbf{0}$.

Hence, we have the stochastic HJB equation

$$V^*(\boldsymbol{\pi}) = \max_u \sum_x R(x, u)\pi(x) + \frac{\partial V^*(\boldsymbol{\pi})}{\partial \boldsymbol{\pi}} \mathbf{f}(\boldsymbol{\pi}, u) + \frac{\tau}{2} \left(\frac{\partial^2 V^*(\boldsymbol{\pi})}{\partial \boldsymbol{\pi}^2} \mathbf{G}(\boldsymbol{\pi}, u) \mathbf{G}(\boldsymbol{\pi}, u)^\top \right).$$

A.4 Gradient and Hessian of a Neural Network

In this section we derive the gradient and Hessian of the input of a value network using back-propagation.

Consider a H -Layer deep Neural network parametrizing a value function $V(\mathbf{x})$ as

$$\begin{aligned} V(\mathbf{x}) &= \mathbf{W}^H \mathbf{z}^{H-1} + \vartheta^H \\ \mathbf{z}^h &= \sigma(\mathbf{a}^h), \quad h = 0, \dots, H-1 \\ \mathbf{a}^h &= \mathbf{W}^h \mathbf{z}^{h-1} + \vartheta^h, \quad h = 1, \dots, H-1 \\ \mathbf{a}^0 &= \mathbf{W}^0 \mathbf{x} + \vartheta^0, \end{aligned}$$

with input \mathbf{x} , weights \mathbf{W}^h , biases ϑ^h and activation function σ . Component wise, the equations are given by

$$\begin{aligned} z_k^h &= \sigma(a_k^h), \quad h = 0, \dots, H-1 \\ a_k^h &= \sum_n w_{kn}^h z_n^{h-1} + \vartheta_k^h, \quad h = 1, \dots, H-1 \\ a_k^0 &= \sum_n w_{kn}^0 x_n + \vartheta_k^0. \end{aligned}$$

A.4.1 The Gradient Computation

Next we calculate the gradient w.r.t. the input i.e., $\partial_{\mathbf{x}} V(\mathbf{x})$. The component wise calculation yield

$$\begin{aligned}\partial_{x_i} V(\mathbf{x}) &= \partial_{x_i} \sum_n w_n^H z_n^{H-1} + \vartheta^H \\ &= \sum_n w_n^H \partial_{x_i} z_n^{H-1}.\end{aligned}$$

where we compute

$$\partial_{x_i} z_k^h = \sigma'(a_k^h) \sum_n w_{kn}^h \partial_{x_i} z_n^{h-1}$$

and σ' denotes the derivative of the activation function σ . For the first layer we compute

$$\partial_{x_i} z_k^0 = \sigma'(a_k^0) w_{ki}^0$$

We define the messages for the partial derivative as $m_{ki}^h := \partial_{x_i} z_k^h$. Therefore, we calculate the message passing as

$$m_{ki}^h = \sigma'(a_k^h) \sum_n w_{kn}^h m_{ni}^{h-1}$$

and

$$m_{ki}^0 = \sigma'(a_k^0) w_{ki}^0$$

Thus, the final equations for the backpropagation of the gradient are given by

$$\begin{aligned}m_{ki}^0 &= \sigma'(a_k^0) w_{ki}^0 \\ \text{For } h &= 1, \dots, H-1 \\ m_{ki}^h &= \sigma'(a_k^h) \sum_n w_{kn}^h m_{ni}^{h-1} \\ \partial_{x_i} V(\mathbf{x}) &= \sum_n w_n^H m_{ni}^{H-1}.\end{aligned}$$

A.4.2 The Hessian Computation

For the Hessian we compute

$$\begin{aligned}\partial_{x_j} \partial_{x_i} V(\mathbf{x}) &= \partial_{x_j} \left\{ \sum_n w_n^H m_{ni}^{H-1} \right\} \\ &= \sum_n w_n^H \partial_{x_j} m_{ni}^{H-1}.\end{aligned}$$

The partial derivatives of the messages are

$$\partial_{x_j} m_{ki}^h = \sigma''(a_k^h) \left(\sum_n w_{kn}^h m_{ni}^{h-1} \right) \left(\sum_n w_{kn}^h m_{nj}^{h-1} \right) + \sigma'(a_k^h) \sum_n w_{kn}^h \partial_{x_j} m_{ni}^{h-1},$$

where σ'' denotes the second derivative of the activation function σ . For the first layer we compute

$$\partial_{x_j} m_{ki}^0 = \sigma''(a_k^0) w_{ki}^0 w_{kj}^0.$$

Next, we define the messages for the second order partial derivative as $\tilde{m}_{kij}^h := \partial_{x_j} m_{ki}^h = \partial_{x_j} \partial_{x_i} z_k^h$. Therefore, we calculate the message passing as

$$\tilde{m}_{kij}^h = \sigma''(a_k^h) \left(\sum_n w_{kn}^h m_{ni}^{h-1} \right) \left(\sum_n w_{kn}^h m_{nj}^{h-1} \right) + \sigma'(a_k^h) \sum_n w_{kn}^h \tilde{m}_{nij}^{h-1}$$

and

$$\tilde{m}_{kij}^0 = \sigma''(a_k^0) w_{ki}^0 w_{kj}^0.$$

Finally, the equations for the back-propagation of the Hessian are given by

$$\begin{aligned}
\tilde{m}_{kij}^0 &= \sigma''(a_k^0) w_{ki}^0 w_{kj}^0 \\
\text{For } h &= 1, \dots, H-1 \\
\tilde{m}_{kij}^h &= \sigma''(a_k^h) \left(\sum_n w_{kn}^h m_{ni}^{h-1} \right) \left(\sum_n w_{kn}^h m_{nj}^{h-1} \right) + \sigma'(a_k^h) \sum_n w_{kn}^h \tilde{m}_{nij}^{h-1} \\
\partial_{x_j} \partial_{x_i} V(\mathbf{x}) &= \sum_n w_n^H \tilde{m}_{nij}^{H-1}.
\end{aligned}$$

B Algorithms

B.1 The Thinning Algorithm

This section contains the thinning algorithm [27] to sample from the presented continuous-time POMDP model, see Algorithm 1.

input : t : time point
 x : state of the latent CTMC, i.e., $X(t) = x$
 $q_{\max} = \max_{u \in \mathcal{U}} \Lambda(x | u)$: maximum exit rate
output : ξ : Waiting time between the start time t and the next jump of $X(t)$

Set current time to start time $T = t$
while *True* **do**
 Sample uniform variable for inverse CDF method $s \sim \text{Uniform}(s | 0, 1)$
 Calculate minimum waiting time sample $\xi_{\min} = \frac{\log u}{q_{\max}}$
 Update current time $T = T + \xi_{\min}$
 Update the filtering dstribution up to T , i.e., $\pi_{[t, T]}$
 Draw $s' \sim \text{Uniform}(s' | 0, 1)$
 if $s' \leq \Lambda(x | \mu(\pi(T)))$ **then**
 Calculate the waiting time $\xi = T - t$
 return *Waiting time ξ and filtering distribution $\pi_{[t, t+\xi]}$*
 end
end

Algorithm 1: Thinning algorithm

B.2 The Collocation Algorithm

In this section we present the collocation algorithm used for the experiments. As a base distribution we use in the experiments e.g., a Dirichlet distribution $\pi^{(i)} \sim \text{Dir}(\pi | \alpha)$. In Algorithm 2 the collocation algorithm can be found for the continuous discrete filter and a finite discrete observation space.

C Experimental Tasks

C.1 Tiger Problem

In the tiger problem [9], the agent has to choose between two doors to open. Behind one door there is a dangerous tiger waiting to eat the agent, thus the agent is supposed to choose the other door to become free. Besides opening a door, the agent can also decide to wait and listen in order to localize the tiger. We adapted the problem to continuous time by defining the POMDP as follows: The state space \mathcal{X} consists of the possible positions of the tiger, *tiger left* and *tiger right*. Executable actions of the agent are *listen*, *open left*, and *open right*. The tiger always stays at the same position therefore all the transition rates $\Lambda(x', x | u)$ are set to zero. When executing an *open* action, the agent receives reward with rate 0.1 for the door without tiger and a negative reward of -1.0 for the door

input : N : Number of collocation samples
 $p(\boldsymbol{\pi})$: Base distribution for collocation samples
 $V_\phi(\boldsymbol{\pi})$: Function approximator for state value function with parameters ϕ
 $\bar{A}_\psi(\boldsymbol{\pi}, u)$: Function approximator for reparametrized advantage function with parameters ψ
output : $\{\hat{\phi}, \hat{\psi}\}$: parameters that have been fitted to approximately solve the HJB equation for $V_{\hat{\phi}}(\boldsymbol{\pi})$ and $\bar{A}_{\hat{\psi}}(\boldsymbol{\pi}, u)$, respectively.

for $i = 1$ **to** N **do**
 Sample collocation beliefs $\boldsymbol{\pi}^{(i)} \sim p(\boldsymbol{\pi})$
 for $u = 1$ **to** N_u **do**
 Compute reset condition $\pi_+^{(i)}(x, y, u) = \frac{p(y|x, u)\pi^{(i)}(x)}{\sum_{x'} p(y|x', u)\pi^{(i)}(x')}, \forall x \in \mathcal{X}, y \in \mathcal{Y}$
 Compute Advantage values for the samples

$$A_\phi(i, u) = \sum_x R(x, u)\pi^{(i)}(x) + V_\phi(\boldsymbol{\pi}^{(i)}) + \tau \sum_{x, x'} \frac{\partial V_\phi(\boldsymbol{\pi}^{(i)})}{\partial \pi(x)} \Lambda(x', x | u)\pi^{(i)}(x') + \tau \lambda \left(\sum_x \pi^{(i)}(x) \sum_y V_\phi(\boldsymbol{\pi}_+^{(i)}(y, u)) - V_\phi(\boldsymbol{\pi}^{(i)}) \right),$$

 where $\{\pi_+^{(i)}(x, y, u)\}_{x \in \mathcal{X}}$ are the components of $\boldsymbol{\pi}_+^{(i)}(y, u)$.
 end
 Compute best action $u^i = \arg \max_u A_\phi(i, u)$
end
Estimate parameters by solving

$$\hat{\phi} = \arg \min_{\phi} \sum_{i=1}^n (A_\phi(i, u^i))^2$$

Recompute all advantage values $A_{\hat{\phi}}(i, u)$ with fitted $\hat{\phi}$ and solve

$$\hat{\psi} = \arg \min_{\psi} \sum_{i=1}^n \sum_{u=1}^{N_u} (\bar{A}_\psi(\boldsymbol{\pi}^{(i)}, u) - \max_{u'} \bar{A}_\psi(\boldsymbol{\pi}^{(i)}, u') - A_{\hat{\phi}}(i, u))^2.$$

Algorithm 2: Collocation algorithm

with tiger. For executing the hearing action, the agent accumulates rewards of rate -0.01 but receives observations with a rate of 2, by hearing the tiger either on the left side (*hear tiger left*) or on the right side (*hear tiger right*). The received information is correct with probability 0.85 thus with probability 0.15 one hears the tiger at the opposite side. The discount factor τ is set to 0.9.

C.2 Slotted Aloha Problem

The slotted aloha transition problem was introduced as POMDP in [8] and deals with decentralized control of stations transmitting packages in a single channel network. The task of the problem is to adjust the sending rate of the stations while the only information about the past transmission state is received at random time points. The state space consists of the number of stations that have a package ready for sending and the past transmission state as observation thus the state space is given by $\mathcal{X} = \{0, \dots, n\} \times \{\text{idle}, \text{transmission}, \text{collision}\}$, where $n = 9$ was chosen to limit the number of states to 30. For our continuous-time POMDP model, we consider a continuous-time adaptation of [8]: Observations arrive with rate 0.5 and contain the past transmission state of the system which contained in the current state, thus $\mathcal{Y} = \{\text{idle}, \text{transmission}, \text{collision}\}$. While the maximum number of packages is not reached, new packages arrive with rate 0.5, leaving the past transmission state unchanged. The stations can send a package simultaneously with rate 5 but do not need to. The action ρ represents the probability with which a package is actually send by a station, resulting in an actual send rate of $5 * \rho/n$.




Figure 4: Advantage updating method applied to slotted aloha problem.

The probability for transmission states given that n packages are available are calculated as

$$\begin{aligned}
 p(\text{idle} \mid \rho) &= (1 - \rho)^n \\
 p(\text{transmission} \mid \rho) &= \rho(1 - \rho)^{n-1} \\
 p(\text{collision} \mid \rho) &= 1 - p(\text{idle} \mid \rho) - p(\text{transmission} \mid \rho)
 \end{aligned}$$

In case of perfect information, the optimal probability for transmission ρ^* can be calculated by maximizing the probability of successful transmission which can be easily obtained by setting the derivative to zero resulting into

$$\rho^* = \arg \max_{\rho} \rho(1 - \rho)^{n-1} = \frac{1}{n}$$

This motivates the discretization of the actions resulting in $\mathcal{U} = \{\frac{1}{n}\}_{n=1,\dots,9}$.

The result of the advantage updating method, analogue to the one presented in the main text is depicted in Fig. 4.

C.3 Gridworld Problem

We consider an agent moving in a 6×6 gridworld with a goal at position $(3, 2)$. There are four actions $\mathcal{U} = \{\text{up}, \text{down}, \text{left}, \text{right}\}$ indicating the direction the agent wants to move next. In our continuous-time setting, the agent moves at an exponentially distributed amount of time with rate of 10. With a probability of 0.7, it moves into the indicated direction, but with probability 0.1 moves into one of the other three directions due to slipping. Movement to invalid fields such as walls is not possible thus for those fields the transition probability is set to zero and the remaining probabilities are renormalized. Being at the goal position provides the agent with a reward with rate 1, otherwise no reward is accumulated. The agent receives noisy signal about his current position at a rate of 2. The signal indicates a field which sampled from a discretized 2D Gaussian distribution with standard deviation of 0.1 centered at the position of the agent.

The result of the collocation method, analogue to the one presented in the main text is depicted in Fig. 5.

D Hyper-Parameters

D.1 Global Hyper-Parameters

Throughout the experiments, we use the following hyper-parameters:




Figure 5: Collocation method applied to gridworld problem.

- The neural networks are parametrized as:

```
linear1 = nn.Linear(in_dim, in_dim)
linear2 = nn.Linear(in_dim, in_dim)
linear3 = nn.Linear(in_dim, out_dim)
```

```
h1 = sigmoid(linear1(x))
h2 = sigmoid(linear2(h1))
out = linear3(h2)
```

- The advantage function network uses

```
in_dim = |X|
out_dim = |U|
```

- The value function network uses

```
in_dim = |X|
out_dim = 1
```

- We use the Adam optimizer with learning rate $\alpha = 10^{-3}$.
- For training we use mini-batches of size $N_b = 256$.

For the collocation method we use:

- $N_o = 1$ optimization step for each mini-batch.
- A training for $N_e = 10000$ episodes.
- A discount decay by scheduling the decay parameter τ for the first 500 steps of optimization, see e.g., [45].

The advantage updating method uses:

- $N_s = 20$ optimization steps per episode, where a mini-batch is sampled from the replay buffer and one optimization step is carried out per mini batch.
- A training for $N_e = 1000$ episodes.
- An exploration process with a damping factor $\kappa = 7.5$ and a decaying noise variance $\sigma \in [1.5, .5]$. For the initial perturbation $\epsilon(u, t = 0)$, we sample the initial condition from the stationary distribution of the Ornstein-Uhlenbeck process.

D.2 Experiment Dependent Hyper-Parameters

For the episode length in the advantage updating algorithm we use:

- Tiger Problem: $l_e = 10$.
- Slotted Aloha Problem: $l_e = 20$.
- Gridworld Problem: $l_e = 5$.

The initial belief for the problems is sampled from the following distributions:

- Tiger Problem: Uniform from the interval $[0.0, 1.0)$.
- Slotted Aloha Problem: Dirichlet distribution with concentration parameter 0.1 for each state dimension.
- Gridworld Problem: Dirichlet distribution with concentration parameter 0.1 for each state dimension.

We sub-sample the episodes for the advantage updating algorithm using the following number of samples:

- Tiger Problem: $N_s = 1000$.
- Slotted Aloha Problem: $N_s = 1000$.
- Gridworld Problem: $N_s = 100$.