

Approximate Q-Learning for Controlled Diffusion Processes and its Near Optimality*

Erhan Bayraktar and Ali Devran Kara [†]

March 10, 2023

Abstract

We study a Q learning algorithm for continuous time stochastic control problems. The proposed algorithm uses the sampled state process by discretizing the state and control action spaces under piece-wise constant control processes. We show that the algorithm converges to the optimality equation of a finite Markov decision process (MDP). Using this MDP model, we provide an upper bound for the approximation error for the optimal value function of the continuous time control problem. Furthermore, we present provable upper-bounds for the performance loss of the learned control process compared to the optimal admissible control process of the original problem. The provided error upper-bounds are functions of the time and space discretization parameters, and they reveal the effect of different levels of the approximation: (i) approximation of the continuous time control problem by an MDP, (ii) use of piece-wise constant control processes, (iii) space discretization. Finally, we state a time complexity bound for the proposed algorithm as a function of the time and space discretization parameters.

Contents

| | | |
|----------|--|----------|
| 1 | Introduction | 2 |
| 1.1 | Preliminaries | 2 |
| 1.2 | Problem Formulation and a Proposed Algorithm | 3 |
| 1.3 | Literature Review | 6 |
| 1.4 | Outline of the Paper | 9 |

*E. Bayraktar is partially supported by the National Science Foundation under grant DMS-2106556 and by the Susan M. Smith chair.

[†]The authors are with the Department of Mathematics, University of Michigan, Ann Arbor, MI, USA, Email: {erhan,alikara}@umich.edu

| | |
|---|-----------|
| 2 Approximate Markov Chains | 10 |
| 2.1 A Markov Chain Construction with Exact Approximation of the Diffusion Process | 10 |
| 2.2 Finite State MDP Construction by Discretization of the State Space | 11 |
| 3 Approximation of the Diffusion Process by a Finite MDP | 13 |
| 3.1 Analysis of term (14) | 13 |
| 3.2 Analysis of term (15) | 14 |
| 3.3 Analysis of term (16) | 16 |
| 3.4 Near Optimality of the Approximate Control | 17 |
| 4 Convergence of the Learning Algorithm and Near Optimality of the Learned Policies | 18 |
| 5 A Discussion on the Convergence Rate of the Learning and the Effect of Sampling and Quantization Rates on the Learning Speed | 23 |
| 6 Conclusion | 26 |
| A Proof of Proposition 1 | 28 |
| B Proof of Proposition 4 | 29 |
| C Proof of Lemma 1 | 32 |
| D Proof of Lemma 2 | 33 |

1 Introduction

The goal of this paper is to develop a convergent learning algorithm for controlled diffusion processes when the decision maker only has access to the state process and the cost realizations, and establish error bounds for the performance of the learned control process compared to the optimal admissible control as a function of the algorithm parameters.

1.1 Preliminaries

We start with the setup of the paper by defining the dynamics of the control problem. The objective is to study a controlled diffusion process, $X(\cdot)$, given

by the following stochastic differential equation

$$X(t) = X_0 + \int_0^t b(X(s), u(s))ds + \int_0^t \sigma(X(s), u(s))dB(s) \quad (1)$$

for $t \geq 0$ where $X(t) \in \mathbb{X}$, and $\mathbb{X} \subset \mathbb{R}$. $B(\cdot)$ is the driving noise which is assumed to be a Wiener process, and $u(\cdot) \in \mathbb{U}$ is the control process with measurable paths. We assume that the control is non-anticipative such that for $0 \leq s < t$, the noise increments $B(t) - B(s)$, are independent of $B(y), u(y)$ for $y \leq s$. We note that the results presented in this paper will be applicable for multidimensional spaces as well, however, we will assume that everything is one-dimensional for ease of notation.

We will later put assumptions on the model, which will guarantee the existence of strong solutions to the diffusion process (1), under admissible control processes (Assumption 1).

The objective of the controller is to minimize the following infinite horizon discounted cost function

$$W_\beta(x_0, u) = E \left[\int_0^\infty e^{-\beta s} c(X(s), u(s)) ds \right] \quad (2)$$

where the expectation is with respect to the initial point $X(0) = x_0$ and the given control process u , the stage-wise cost function is denoted by $c : \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{R}$, and β is some discount factor.

The optimal cost is denoted by

$$W_\beta^*(x_0) := \inf_u W_\beta(x_0, u)$$

where the infimum is taken over all admissible control processes.

1.2 Problem Formulation and a Proposed Algorithm

Our goal is to provide a reinforcement learning algorithm which converges to an approximately optimal control policy under certain assumptions. In particular, we will show that if the standard Q learning algorithm ([47, 41, 13]) is used for the state process after time and space discretization, we can find sufficient conditions for the algorithm to converge. Furthermore, we will provide provable error bounds for the performance of the policies learned through these iterations. We note that the standard Q learning algorithm is not readily applicable for the diffusion model we consider in the paper. The classical Q learning algorithm is designed for finite Markov chains, where the Markov property of the data is crucially used to prove

the convergence of the algorithm. In the model we consider here, time discretization breaks the Markov property of the induced chain unless the control policies are selected carefully, namely selected as piece-wise constant functions. Furthermore, even when the sampled process which lives on the original continuous state space is Markov, the process constructed using space discretization will not be a Markov chain. Hence, one needs to alter the convergence proof significantly when the space is discretized, in order to show the classical Q learning algorithm converges under time and space discretization.

We assume that the drift, diffusion, and the cost functions b , σ and c are not known. By observing the state process $X(t)$ and the cost realizations $c(X(t), u(t))$, we try to learn the Q values of a finite Markov decision process (MDP), which will be shown to well approximate the original problem. Furthermore, we keep track of the state and the cost processes via sampling the time and discretizing \mathbb{X} and \mathbb{U} spaces. To this end, we fix a sampling interval $h > 0$, and finite subsets $\mathbb{X}_h \subset \mathbb{X}$, $\mathbb{U}_h \subset \mathbb{U}$. We put the h dependence on the finite spaces to recover the cases where the space discretization rate depends on the time discretization rate. We also define the mapping $\phi_{\mathbb{X}} : \mathbb{X} \rightarrow \mathbb{X}_h$, to map the original value of the state variables to the discrete set \mathbb{X}_h , e.g. a natural choice would be a nearest neighbour map.

For the exploration phase, we use a piecewise constant control process, such that at $i \times h$, for $i = 0, 1, \dots$, some $\hat{u} \in \mathbb{U}_h$ is chosen according to the exploration policy and is applied for the time interval $[i \times h, (i + 1) \times h]$. Hence, the resulting exploration process $\hat{u}(t)$ is a piecewise constant process.

After we fix the exploration control process $\hat{u}(t)$, we observe the controlled Markov chain process X_n

$$X_n := X(n \times h), \quad \text{for } n = 0, 1, 2, \dots$$

Note that X_n takes values from a continuous set \mathbb{X} . We further map these values to the finite set \mathbb{X}_h , to construct the stochastic process \hat{X}_n , which is defined as

$$\hat{X}_n := \phi_{\mathbb{X}}(X_n) \tag{3}$$

The process \hat{X}_n is the discretized version of the sampled diffusion process, and hence it takes values from \mathbb{X}_h . We note that the learning algorithm will be constructed using the process $\{\hat{X}_n\}_n$.

We now present the proposed algorithm formally as follows:

1. Chose a sampling interval $h > 0$.

2. Choose finite subsets $\mathbb{X}_h \subset \mathbb{X}$ and $\mathbb{U}_h \subset \mathbb{U}$.
3. Define a mapping $\phi_{\mathbb{X}} : \mathbb{X} \rightarrow \mathbb{X}_h$ (e.g. a nearest neighbour map).
4. Select a \mathbb{U}_h valued piecewise constant exploration process $\hat{u}(t)$.
5. Observe the discrete process \hat{X}_n (see (3)), and for all $(\hat{x}, \hat{u}) \in \mathbb{X}_h \times \mathbb{U}_h$, initiate the Q values by choosing $Q_0(\hat{x}, \hat{u})$, choose the learning rates $\alpha_k(\hat{x}, \hat{u})$ (see Assumption 2), using the cost realization $c(X(n \times h), \hat{u})$ update

$$Q_{k+1}(\hat{x}, \hat{u}) = (1 - \alpha_k(\hat{x}, \hat{u}))Q_k(\hat{x}, \hat{u}) + \alpha_k(\hat{x}, \hat{u}) \left(c(X(k \times h), \hat{u}) \times h + \beta_h \min_{v \in \mathbb{U}_h} Q_k(\hat{X}_{k+1}, v) \right) \quad (4)$$

where $\beta_h = e^{-\beta \times h}$, and \hat{X}_{k+1} is the sampled state we observe following $\hat{X}_k = \hat{x}$.

6. In Section 4, we will prove that under suitable assumptions, the iterations $Q_k : \mathbb{X}_h \times \mathbb{U}_h \rightarrow \mathbb{R}$, converge almost surely to some $Q^* : \mathbb{X}_h \times \mathbb{U}_h \rightarrow \mathbb{R}$. In Section 2, we will show that limit values are the Q -values of some finite controlled Markov chain. Using Q^* , we define a mapping $\gamma_h : \mathbb{X}_h \rightarrow \mathbb{U}_h$ such that for any $\hat{x} \in \mathbb{X}_h$

$$\gamma_h(\hat{x}) = \arg \min_{\hat{u} \in \mathbb{U}_h} Q^*(\hat{x}, \hat{u}).$$

7. Define the control process $u_h(t)$ such that

$$u_h(t) = \gamma_h(\phi_{\mathbb{X}}(X(i \times h)), \text{ for } t \in [i \times h, (i+1) \times h],$$

that is u_h is a piece-wise constant process, which changes value at the sampling instances according to the learned map γ_h .

We are then interested in the following question:

Problem: If we use the learned control u_h for the continuous time model, what do we lose compared the optimal admissible control process. In other words, we are interested in the difference

$$W_{\beta}(x_0, u_h) - W_{\beta}^*(x_0). \quad (5)$$

We will try to bound this term in terms of the sampling interval h and the state and control action spaces' discretization errors $L_{\mathbb{X}}$ and $L_{\mathbb{U}}$ that are defined as

$$L_{\mathbb{X}} := \sup_{x \in \mathbb{X}} |\phi_{\mathbb{X}}(x)|,$$

$$L_{\mathbb{U}} := \sup_{u \in \mathbb{U}} \min_{u_h \in \mathbb{U}_h} |u - u_h|.$$

We will also analyze the asymptotic case, that is, we will show that under certain assumptions

$$W_{\beta}(x_0, u_h) - W_{\beta}^*(x_0) \rightarrow 0$$

as $h, L_{\mathbb{X}}, L_{\mathbb{U}} \rightarrow 0$.

To answer these questions, we will show that the iterations in (4) converge to the optimal Q values of some finite Markov chain which approximates the value function of the diffusion process in (1) with the cost function (2).

1.3 Literature Review

The results and the techniques used in the paper are related to reinforcement learning algorithms for continuous and discrete time control problems, finite time-space approximations of diffusion processes, and approximation methods for Markov decision processes. In the following, we summarize the related works in the literature by their main objectives:

Reinforcement learning for discrete-time stochastic control problems Optimal control of stochastic (or deterministic) dynamical systems typically requires a perfect knowledge about the system components, however, the correct underlying model of the control problem is usually unknown or too complicated to work with. The objective of reinforcement learning is to estimate or learn the relevant information about the problem such as the value function or the optimal policy by interacting and observing the system. Majority of the literature, however, deals with discrete-time and sequential decision making problems. A popular reinforcement learning algorithm, called Q learning ([47]), for example, is proposed to learn the Q factors, which is closely related to the value function, for discrete time and space problems. Q learning is guaranteed to converge under mild assumptions for discrete settings without any information on the transition models or the cost function of the problem by observing the state and cost realizations ([41, 13]).

Even though Q learning is simple to implement and requires minimal knowledge about the system, it is not effective for large or continuous state and action spaces. To overcome the dimension challenges, one can try to learn an approximate version of the problem. The approximation can be done through various methods: One can use function approximations for the optimal value function (see [37, 42]). To approximate the value functions, neural networks, state aggregation, or linear approximation techniques with finitely many linearly independent basis functions can be used. For state aggregation and linear approximation methods, convergence can be shown, however, a rigorous error analysis is usually not available. Some related work includes [36, 28, 11, 38, 35] and references therein. For the aforementioned works, typically, either a careful parametrization of the value functions or strong density assumptions on the transition kernels are required. However, for general continuous time problems, these assumptions might be too strong. For example, the sampled continuous time stochastic controlled process we will study in this paper, can only be shown to have weakly continuous transition dynamics (see Lemma 1 and Lemma 2). In [15], it is shown that, by choosing a finite subset of the action space and discretizing the state space, one can consistently learn nearly optimal control policies for systems with weakly continuous transition models.

Approximations for continuous time control problems Space approximation methods allow one to use learning algorithms for control problems with large state spaces, however, these methods only work for discrete time decision making problems. Control problems in continuous time, in general, are not feasible to work with numerically. Due to complex nature of the problems, solutions of the control process may not be available, which results in approximation attempts on the time domain, either through the state process or the optimality equation. [22] provides a general Markov chain approximation method for the controlled diffusion processes, by directly approximating the continuous time state process with a finite space controlled Markov chain. It is shown that under certain 'consistency' conditions, state process and value function approximations are asymptotically close to the solutions of the continuous time process. [19, 20] establish convergence rates for such approximation methods. Improved convergence rates for finite approximation methods are later presented in [5, 4]. We note that, these works study the approximation of an available model and do not focus on learning an approximate model when the dynamics are not available to the decision maker.

Reinforcement learning in continuous time Learning and planning in continuous time problems become much more challenging in continuous

time mainly because of the complex dynamics and optimality equation of the problems.

We first note that our main objective and contribution is to rigorously prove convergence of the discretization based model-free learning algorithms for general diffusion processes and to provide provable error bounds that clearly indicate the effect of space and time discretization.

A large number of papers dealing with learning of continuous time control problems (learning of value functions or control policies), considers linear dynamics and cost setting, and develop algorithms and theoretical results using the structural properties of this limited setting, see e.g. [32, 8, 43, 34, 18, 45, 46] for some of the related papers that work with linear models. In our paper, we consider general non-linear controlled diffusion dynamics, where the only assumption we put on the dynamics is a standard continuity assumption, which is required for existence results.

Another large set of studies in learning of continuous time control problems consider deterministic dynamics see e.g. [40, 48, 27, 18, 23, 50, 31]. We note that the approximation, convergence and structural analysis are considerably more challenging for continuous time *stochastic* control problems. We further note that some of these studies focus on approximate solutions to the optimal control problem with an available model, instead of considering model-free learning methods.

Another related direction includes estimation and learning of dynamics for stochastic differential equations see e.g. [49, 6, 33, 26]. In these papers, the focus is on the learning of models for control-free stochastic processes governed by stochastic differential equations (SDEs), which differs from our objective of learning optimal value functions and-or near optimal policies in controlled stochastic processes.

Learning methods are also used to find (approximate) solutions to partial differential equations and to solve HJB equations (approximately) see e.g. [40, 1, 12, 24]. Note that these works approach to the solution problem with a model in hand (i.e. model-based not model-free), where the focus is on finding the solution of the PDEs and HJB equations which is by itself a challenging problem.

The closest papers to ours are [39, 3, 30]. [30] proposes a learning method using time and space discretizations for a general controlled diffusion process, and proves that the learned approximations are asymptotically optimal as the discretization parameters converge to 0. In our paper, we present error bounds in terms of the space and time discretization rate, which explicitly shows the effects of discretization. [39, 3] proposes ‘advantage functions’ for continuous time learning algorithms building on the observation that Q

values are no longer informative for continuous time settings, and this observation is also widely used in the literature. In our results, we show that this claim should be approached with care. We show that as long as the space discretization rate is at least as high as the time discretization rate, Q values of the approximate MDP model well approximate the value function of the controlled diffusion problem.

Our contributions In this paper, we will study an approximate Q learning algorithm for general continuous time stochastic control problems by discretizing the time, and state and action spaces. The discretization in time will be assumed to be uniform, however, state and action spaces can be discretized in a non-uniform way. We will show that the algorithm will converge under mild assumptions, even when the space quantization is non-uniform. Different from earlier works, we will then provide error bounds and convergence rates in terms of the discretization parameters. The bounds will suggest that, even though the algorithm converges for general space quantization, the performance of the learned value functions and policies will depend on the quantization scheme. Furthermore, provided error bounds will emphasize the effects of different levels of the approximation such as estimating the diffusion process with a controlled Markov chain, use of piece-wise constant policies, and state and action space quantization. Lastly, we will also discuss the effect of discretization parameters on the learning speed of the algorithm.

1.4 Outline of the Paper

In Section 2, we construct a finite space Markov decision process (MDP) that will serve to approximate the diffusion process (1). In particular, in Section 2.1, we present a controlled discrete time control process, which has the same distributions as the diffusion process at the sampling instances, when the diffusion process is controlled using piece-wise constant control functions. In Section 2.2, we present a finite space controlled Markov chain, that is constructed based on the MDP from Section 2.1, using state aggregation methods.

In Section 3, we analyze the differences between the value function of the diffusion process (1), and the value function of the finite space MDP constructed in Section 2.2. Furthermore, we provide upper bounds for the error (or regret) of the control policy designed for the finite space MDP, when it is used to control the diffusion process, in terms of the approximation parameters; where the comparison is with respect to the performance of the optimal admissible control process.

In Section 4, we show that the algorithm (4) converges, under certain assumptions, to the Q values of the finite space MDP constructed in Section 2.2, and thus, we provide error bounds for the learned policy using the results from Section 3.

Finally, in Section 5, towards a practical purpose, we analyze the convergence rate of the algorithm (4), with respect to the sampling interval lengths, h , and the number of iterations where we provide time complexity bounds.

2 Approximate Markov Chains

In this section, we provide two controlled Markov chains which will help us to analyze the error term (5).

2.1 A Markov Chain Construction with Exact Approximation of the Diffusion Process

The first Markov chain we will present will have the same finite dimensional distributions with the sampled diffusion process under piece-wise constant control processes.

Let \mathbb{X} be the state space, and \mathbb{U}_h be the control action space of the Markov chain.

We define the transition probabilities as follows: For any $k \in \mathbb{Z}_+$, distribution of the state X_k conditioned on the past state and action variables, is determined by the diffusion process (1), such that, conditioned on $(x_{k-1}, \dots, x_0, u_{k-1}, \dots, u_0)$, X_k has the same distribution as

$$X(h) = x_{k-1} + \int_0^h b(X(s), u_{k-1}) ds + \int_0^h \sigma(X(s), u_{k-1}) dB(s). \quad (6)$$

Hence, for any $A \in \mathcal{B}(\mathbb{X})$

$$Pr(X_k \in A | x_{[0, k-1]}, u_{[0, k-1]}) = \mathcal{T}_h(A | x_{k-1}, u_{k-1})$$

where $(x, u)_{[0, k-1]} := x_0, \dots, x_{k-1}, u_0, \dots, u_{k-1}$, such that

$$\mathcal{T}_h(dx_k | x_{k-1}, u_{k-1}) \sim X(h)$$

where $X(h)$ determined by (6) and where \mathcal{T}_h is the transition kernel of the Markov chain which is a stochastic kernel from $\mathbb{X} \times \mathbb{U}_h$ to \mathbb{X} .

We also define a stage-wise cost function $c_h : \mathbb{X} \times \mathbb{U}_h \rightarrow \mathbb{R}_+$ such that for any $(x, u) \in \mathbb{X} \times \mathbb{U}_h$

$$c_h(x, u) := c(x, u) \times h$$

where c is the cost function of the diffusion process (see (2)).

We now define the infinite horizon discounted cost function

$$J_{\beta_h}(x_0, \gamma) := E_{x_0}^{\mathcal{T}_h, \gamma} \left[\sum_{k=0}^{\infty} \beta_h^k c_h(X_k, U_k) \right] \quad (7)$$

where $\beta_h := e^{-\beta \times h}$, and γ is an *admissible policy*. An *admissible policy* is a sequence of control functions $\{\gamma_k, k \in \mathbb{Z}_+\}$ such that γ_k is measurable with respect to the σ -algebra generated by the information variables $I_k = \{X_{[0,k]}, U_{[0,k-1]}\}$, $k \in \mathbb{N}$, $I_0 = \{Y_0\}$, where

$$U_k = \gamma_k(I_k), \quad k \in \mathbb{Z}_+,$$

are the \mathbb{U}_h -valued control actions. We define Γ to be the set of all such admissible policies. The optimal cost function is defined as

$$J_{\beta_h}^*(x_0) := J_{\beta_h}(x_0, \gamma). \quad (8)$$

Remark 1. *An important property of the MDP model we will make use of, is the following one: suppose that we are given an admissible policy $\gamma \in \Gamma$ defined for the MDP, we define the following control process $u(t)$ such that*

$$u(t) = \gamma(X(k \times h)), \quad \text{for } t \in [k \times h, (k+1) \times h] \quad (9)$$

which is a piece-wise constant control process. Then, the controlled Markov chain state process X_k under the policy γ and the controlled diffusion process $X(t)$ under the control process defined in (9) have the same distributions at the sampling instances if they start from the same initial points, that is, for any $k \in \mathbb{Z}_+$,

$$X_k \sim X(k \times h).$$

2.2 Finite State MDP Construction by Discretization of the State Space

We now construct an MDP with a finite state space by discretizing the state space \mathbb{X} .

We start by choosing a collection of disjoint sets $\{B_i\}_{i=1}^M$ such that $\cup_i B_i = \mathbb{X}$, and $B_i \cap B_j = \emptyset$ for any $i \neq j$. Furthermore, we choose a

representative state, $\hat{x}_i \in B_i$, for each disjoint set. For this setting, we denote the new finite state space by $\mathbb{X}_h := \{\hat{x}_1, \dots, \hat{x}_M\}$. We put the dependence on the parameter h , since we will let the size of the finite set, M go to ∞ as $h \rightarrow 0$. The mapping from the original state space \mathbb{X} to the finite set \mathbb{X}_h is done via

$$\phi_{\mathbb{X}}(x) = \hat{x}_i \quad \text{if } x \in B_i. \quad (10)$$

Furthermore, we choose a weight measure $\pi^*(\cdot) \in \mathcal{P}(\mathbb{X})$ on \mathbb{X} such that $\pi^*(B_i) > 0$ for all $i \in \{1, \dots, M\}$. We now define normalized measures using the weight measure on each separate quantization bin B_i such that

$$\hat{\pi}_{\hat{x}_i}^*(A) := \frac{\pi^*(A)}{\pi^*(B_i)}, \quad \forall A \subset B_i, \quad \forall i \in \{1, \dots, M\} \quad (11)$$

that is $\hat{\pi}_{\hat{x}_i}^*$ is the normalized weight measure on the set B_i , \hat{x}_i belongs to.

We now define the cost and transition kernels for this finite set using the normalized weight measures such that for any $\hat{x}_i, \hat{x}_j \in \mathbb{X}_h$

$$\begin{aligned} C_h^*(\hat{x}_i, u) &= \int_{B_i} c_h(x, u) \hat{\pi}_{\hat{x}_i}^*(dx) \\ P_h^*(\hat{x}_j | \hat{x}_i, u) &= \int_{B_i} \mathcal{T}_h(B_j | x, u) \hat{\pi}_{\hat{x}_i}^*(dx). \end{aligned} \quad (12)$$

where \mathcal{T}_h is the transition model for the MDP constructed in Section 2.1.

Having defined the finite state space \mathbb{X}_h , the cost function C_h^* and the transition model P_h^* , we can now introduce the optimal value function for this finite model. We denote the optimal value function, which is defined on \mathbb{X}_h , by $\hat{J}_{\beta_h} : \mathbb{X}_h \rightarrow \mathbb{R}$. Note that \hat{J}_{β_h} satisfies the following Bellman equation for any $\hat{x}_i \in \mathbb{X}_h$:

$$\hat{J}_{\beta_h}(\hat{x}_i) = \inf_{u \in \mathbb{U}_h} \left\{ C_h^*(\hat{x}_i, u) + \beta_h \sum_{\hat{x}_1 \in \mathbb{X}_h} \hat{J}_{\beta_h}(\hat{x}_1) P_h^*(\hat{x}_1 | \hat{x}_i, u) \right\} \quad (13)$$

We can easily extend this function over the state space \mathbb{X} by making it constant over the quantization bins. In other words, if $\hat{x}_i \in B_i$, where $\cup_i B_i = \mathbb{X}$, for any $x \in B_i$, we write

$$\hat{J}_{\beta_h}(x) := \hat{J}_{\beta_h}(\hat{x}_i).$$

Furthermore, the following equation follows directly from the dynamic programming principle for the Q values of the finite MDP:

$$Q_h^*(\hat{x}_i, u) = C_h^*(\hat{x}_i, u) + \beta_h \sum_{\hat{x}_1} \min_v Q_h^*(\hat{x}_1, v) P_h^*(\hat{x}_1 | \hat{x}_i, u).$$

Remark 2. We will prove that the iterations (4) converge to some Q^* which satisfies the above equation, under suitable conditions on the diffusion process.

We further define uniform error bounds resulting from the discretization of the state and actions spaces such that

$$L_{\mathbb{X}} := \sup_{x \in \mathbb{X}} |x - \phi_{\mathbb{X}}(x)|,$$

$$L_{\mathbb{U}} := \sup_{u \in \mathbb{U}} \min_{u_h \in \mathbb{U}_h} |u - u_h|.$$

We note that later in the paper (see Corollary 1) we will see that, for the approximation error of the diffusion process to go to 0 with increasing discretization rates, we will need $L_{\mathbb{X}}, L_{\mathbb{U}} < h$.

3 Approximation of the Diffusion Process by a Finite MDP

Recall that we are interested in the term

$$W_{\beta}(x_0, u_h) - W_{\beta}^*(x_0)$$

where u_h the control process obtained with γ_h that is learned via (4). The first term represents the cost induced by the application of the control designed for the finite space MDP, when it is used for the continuous time model. The second term represents the optimal cost for the continuous time model. In Section 4, we will prove that the learned policy γ_h is optimal for the MDP defined in Section 2.2. Hence, in this section, we will assume that γ_h solves (13).

We write the following:

$$W_{\beta}(x_0, u_h) - W_{\beta}^*(x_0) = W_{\beta}(x_0, u_h) - J_{\beta_h}(x_0, \gamma_h) \quad (14)$$

$$+ J_{\beta_h}(x_0, \gamma_h) - J_{\beta_h}^*(x_0) \quad (15)$$

$$+ J_{\beta_h}^*(x_0) - W_{\beta}^*(x_0). \quad (16)$$

In what follows, we will analyze each term separately.

3.1 Analysis of term (14)

The following result provides a bound for the difference between the value function of a diffusion process controlled with piece-wise constant control

processes and the value function of an MDP controlled with an admissible policy.

Assumption 1. *For the diffusion process given in (1), we assume that*

- $\sup_{x,u} |b(x,u)| \leq B$, and $\sup_{x,u} |\sigma(x,u)| \leq B$, for some $B < \infty$.
- $|b(x,u) - b(x',u')| + |\sigma(x,u) - \sigma(x',u')| \leq K(|x-x'| + |u-u'|)$ for some $K < \infty$ and for any $x, x' \in \mathbb{X}$, and $u, u' \in \mathbb{U}$.
- $\sup_{x,u} |c(x,u)| \leq C$ for some $C < \infty$.
- $|c(x,u) - c(x',u')| \leq K(|x-x'| + |u-u'|)$ for some $K < \infty$ and for any $x, x' \in \mathbb{X}$, and $u, u' \in \mathbb{U}$.
- The process is nondegenerate such that $\sigma^2(x,u) > 0$ for every $x, u \in \mathbb{X} \times \mathbb{U}$.

Note that these assumptions are sufficient for the existence of a unique strong solution to (1) under admissible control processes.

Proposition 1. *Let $\gamma \in \Gamma$ be an admissible policy for the sampled controlled Markov chain, and $u_h(\cdot)$ be the corresponding piece-wise constant control process for the diffusion process such that*

$$u_h(t) = \gamma(x(i \times h)), \text{ for } t \in [i \times h, (i+1) \times h].$$

Under Assumption 1, we have that for any $x_0 \in \mathbb{X}$

$$|W_\beta(x_0, u_h) - J_{\beta_h}(x_0, \gamma)| \leq hB + \frac{KBh}{1 - e^{-\beta h}} \left(h + \sqrt{\frac{2h}{\pi}} \right).$$

Proof. The proof can be found in Appendix A. □

3.2 Analysis of term (15)

For the analysis of (15), we will need to calculate the Lipschitz constants of the sampled controlled Markov chain in terms of the Lipschitz constants of the diffusion process introduced in Assumption 1.

In what follows, we will focus on the controlled Markov chain constructed in Section 2.1. Recall that we have

$$\mathcal{T}_h(dx_k | x_{k-1}, u_{k-1}) \sim X(h)$$

where

$$X(h) = x_{k-1} + \int_0^h b(X(s), u_{k-1}) ds + \int_0^h \sigma(X(s), u_{k-1}) dB(s).$$

Furthermore,

$$\begin{aligned}\beta_h &:= e^{-\beta h} \\ c_h(x, u) &:= c(x, u) \times h.\end{aligned}$$

Lemma 1. *Under Assumption 1,*

$$W_1(\mathcal{T}_h(\cdot|y, u), \mathcal{T}_h(\cdot|x, u)) \leq |x - y| e^{(K + \frac{K^2}{2})h}.$$

where W_1 denotes the first order Wasserstein distance.

Proof. The proof can be found in Appendix C. \square

Lemma 2. *Under Assumption 1, if $h < 1$*

$$W_1(\mathcal{T}_h(\cdot|x, u), \mathcal{T}_h(\cdot|x, \hat{u})) \leq |u - \hat{u}| 2K e^{2K^2 h}.$$

where W_1 denotes the first order Wasserstein distance.

Proof. The proof can be found in Appendix D. \square

Proposition 2. *Under Assumption 1, if $\beta > K + \frac{K^2}{2}$, and if $h < 1$,*

$$\begin{aligned}\sup_{x_0 \in \mathbb{X}} |J_{\beta_h}(x_0, \hat{\gamma}_h) - J_{\beta_h}^*(x_0)| &\leq \frac{Kh - Khe^{h(K + \frac{K^2}{2} - \beta)} + 2K^2he^{h(2K^2 - \beta)}}{(1 - e^{-\beta h})(1 - e^{h(K + \frac{K^2}{2} - \beta)})} L_{\mathbb{U}} \\ &\quad + \frac{Kh}{(1 - e^{-\beta h})(1 - e^{h(K + \frac{K^2}{2} - \beta)})} L_{\mathbb{X}}\end{aligned}$$

Proof. The proof is a direct implication of [15, Theorem 2.6], which states that

$$\sup_{x_0 \in \mathbb{X}} |J_{\beta_h}(x_0, \hat{\gamma}_h) - J_{\beta_h}^*(x_0)| \leq \frac{\alpha_c^{\mathbb{X}} - \beta_h \alpha_{\mathcal{T}}^{\mathbb{X}} \alpha_c^{\mathbb{U}} + \beta_h \alpha_c^{\mathbb{X}} \alpha_{\mathcal{T}}^{\mathbb{U}}}{(1 - \beta)(1 - \beta_h \alpha_{\mathcal{T}}^{\mathbb{X}})} L_{\mathbb{U}} + \frac{\alpha_c^{\mathbb{X}}}{(1 - \beta_h)(1 - \beta_h \alpha_{\mathcal{T}}^{\mathbb{X}})} L_{\mathbb{X}},$$

for constants $\alpha_c^{\mathbb{X}}, \alpha_c^{\mathbb{U}}, \alpha_{\mathcal{T}}^{\mathbb{X}}, \alpha_{\mathcal{T}}^{\mathbb{U}} < \infty$ such that

- $|c(x, u) - c(x', u)| \leq \alpha_c^{\mathbb{X}} |x - x'|,$

- $|c(x, u) - c(x, u')| \leq \alpha_c^U |u - u'|$,
- $W_1(\mathcal{T}(\cdot|x, u), \mathcal{T}(\cdot|x', u)) \leq \alpha_{\mathcal{T}}^X |x - x'|$,
- $W_1(\mathcal{T}(\cdot|x, u), \mathcal{T}(\cdot|x, u')) \leq \alpha_{\mathcal{T}}^U |u - u'|$,

Hence, the result follows from Lemma 1 and Lemma 2, by noting that $\alpha_c^X, \alpha_c^U \leq hK$. \square

3.3 Analysis of term (16)

Recall that (16) deals with the optimal value function of the sampled controlled Markov chain and the optimal value function of the controlled diffusion process. Hence, we make use of finite difference approximation methods for Bellman equations. The following result, taken from [21, 14], gives an upper bound on the performance loss of the piece-wise constant policies applied for the diffusion processes.

Lemma 3 ([21, 14]). *Let u_h^* denote the piece-wise constant control process, which is constant over the intervals $[kh, k(h+1))$ for $k \in \{0, 1, \dots\}$, that achieves the minimum cost for (2) over such piece-wise constant policies. Under Assumption 1, we have*

$$W_\beta(x_0, u_h^*) - W_\beta^*(x_0) \leq Nh^{\frac{1}{4}}$$

for some constant $N < \infty$, which only depends on the discount factor β and the Lipschitz coefficients of the model.

Remark 3. *One might expect that if the dynamics of the diffusion process are changing slowly, and the cost function does not have rapid changes with respect to the state process, then the time discretization leads to smaller performance losses. Indeed, as it is shown in [14, Proposition 2.4], if W_β and c are regular enough, namely if*

$$\sup_u \left(\|L_u(L_u W_\beta)\|_\infty + \|L_u c\|_\infty \right) < \infty$$

where L_u is the generator function of the diffusion process for some control action u , then one might have

$$W_\beta(x_0, u_h^*) - W_\beta^*(x_0) \leq Nh.$$

We are now ready to analyze (16).

Proposition 3. *Under Assumption 1*

$$J_{\beta_h}^*(x_0) - W_{\beta}^*(x_0) \leq hB + \frac{KBh}{1 - e^{-\beta h}} \left(h + \sqrt{\frac{2h}{\pi}} \right) + Nh^{\frac{1}{4}}.$$

Proof. We start by the following bound

$$\begin{aligned} J_{\beta_h}^*(x_0) - W_{\beta}^*(x_0) &\leq J_{\beta_h}^*(x_0) - W_{\beta}(x_0, u_h^*) + W_{\beta}(x_0, u_h^*) - W_{\beta}^*(x_0) \\ &\leq J_{\beta_h}^*(x_0) - W_{\beta}(x_0, u_h^*) + Nh^{\frac{1}{4}} \end{aligned}$$

where the last bound follows directly from Theorem 3. For the first term, let γ_h^* denote the corresponding control policy for the controlled Markov chain (that has the same law at the sampling instances as u_h^*). Note that this policy is not necessarily optimal for the sampled controlled Markov chain. Hence, we have the following bound:

$$\begin{aligned} J_{\beta_h}^*(x_0) - W_{\beta}(x_0, u_h^*) &\leq J_{\beta_h}(x_0, \gamma_h^*) - W_{\beta}(x_0, u_h^*) \\ &\leq hB + \frac{KBh}{1 - e^{-\beta h}} \left(h + \sqrt{\frac{2h}{\pi}} \right). \end{aligned}$$

The last bound follows from Theorem 1.

Combining what we have so far, we can conclude that

$$J_{\beta_h}^*(x_0) - W_{\beta}^*(x_0) \leq hB + \frac{KBh}{1 - e^{-\beta h}} \left(h + \sqrt{\frac{2h}{\pi}} \right) + Nh^{\frac{1}{4}}.$$

□

3.4 Near Optimality of the Approximate Control

Combining the results, we have presented so far, we can now state the main theorem of this section:

Theorem 1. *Under Assumption 1, we have*

$$W_{\beta}(x_0, u_h) - W_{\beta}^*(x_0) \leq e(h, L_{\mathbb{X}}, L_{\mathbb{U}})$$

where

$$\begin{aligned}
e(h, L_{\mathbb{X}}, L_{\mathbb{U}}) := & 2 \left(hB + \frac{KBh}{1 - e^{-\beta h}} \left(h + \sqrt{\frac{2h}{\pi}} \right) \right) \\
& + \frac{Kh - Khe^{h(K + \frac{K^2}{2} - \beta)} + 2K^2he^{h(2K^2 - \beta)}}{(1 - e^{-\beta h})(1 - e^{h(K + \frac{K^2}{2} - \beta)})} L_{\mathbb{U}} \\
& + \frac{Kh}{(1 - e^{-\beta h})(1 - e^{h(K + \frac{K^2}{2} - \beta)})} L_{\mathbb{X}} \\
& + Nh^{\frac{1}{4}}.
\end{aligned} \tag{17}$$

where u_h is the control process obtained with γ_h , which is optimal for the finite MDP model constructed in Section 2.2.

Note that, as stated in Remark 3, the last term in Theorem 1 can be replaced by Nh , if the dynamics and the cost function are regular enough, e.g. if they do not change rapidly with time.

Corollary 1. *For small h , the upper bound derived in Theorem 1, behaves as*

$$C \left(\sqrt{h} + \frac{L_{\mathbb{X}} + L_{\mathbb{U}}}{h} + h^{\frac{1}{4}} \right)$$

for some $C < \infty$. This representation makes the distinction between the effects of different steps of the approximation. The first term results from the Markov chain approximation of the diffusion process, the second term is due to the state and action space discretization, and finally the last term is due to the piece-wise constant control processes. We note again that the last term may be replaced with h if the dynamics are regular enough.

4 Convergence of the Learning Algorithm and Near Optimality of the Learned Policies

In this section we present the main results of the paper.

We first show that the iterations in (4) converge to the optimal Q values of the approximate controlled Markov chain constructed in Section 2.2. Recall that the Q value iterations are in the following form:

$$\begin{aligned}
Q_{k+1}(\hat{x}, \hat{u}) &= (1 - \alpha_k(\hat{x}, \hat{u}))Q_k(\hat{x}, \hat{u}) \\
&\quad + \alpha_k(\hat{x}, \hat{u}) \left(c(X(k \times h), \hat{u}) \times h + \beta_h \min_{v \in \mathbb{U}_h} Q_k(\hat{X}_{k+1}, v) \right)
\end{aligned} \tag{18}$$

where $\beta_h = e^{-\beta \times h}$, and \hat{X}_{k+1} is the discretized sampled state we observe following $\phi_{\mathbb{X}}(x(i \times h)) = \hat{x}$. Furthermore, $\phi_{\mathbb{X}}$ maps the original state space \mathbb{X} to the finite subset \mathbb{X}_h .

Assumption 2.

i. $\alpha_k(\hat{x}, \hat{u}) = 0$ unless $(\hat{X}_k, u_k) = (\hat{x}, \hat{u})$. Furthermore,

$$\alpha_k(\hat{x}, \hat{u}) = \frac{1}{1 + \sum_{t=0}^k \mathbf{1}_{\{\hat{X}_t = \hat{x}, u_t = \hat{u}\}}}.$$

This implies $\alpha_k(\hat{x}, \hat{u}) = \frac{1}{t+1}$ if we have visited (\hat{x}, \hat{u}) pair t many times until time k , i.e. the learning rates are linear.

ii. The controlled diffusion process converges to its unique invariant measure under the exploration policy.

iii. Every $\hat{x} \in \mathbb{X}_h$ and $\hat{u} \in \mathbb{U}_h$ is visited infinitely often during exploration.

Remark 4. For the stability assumption (ii), we need the exploration policy to be a stabilizing policy such that it leads the process to its invariant measure. Since we use piece-wise constant policies for exploration, the stability can be tested using the Lyapunov type stability criteria for the resulting discrete time Markov decision process (see Section 2.1). E.g. let $C \subset \mathbb{X}$ be a compact set, $b \in \mathbb{R}$, $\epsilon > 0$, and $V : \mathbb{X} \rightarrow \mathbb{R}_+$ (e.g. $V(x) = |x|$), if the following is satisfied for all $x \in \mathbb{X}$:

$$\int_{\mathbb{X}} V(y) \mathcal{T}_h^{\gamma}(dy|x) = E[V(x_{t+1})|x_t = x] \leq V(x) - \epsilon + b \mathbf{1}_{\{x \in C\}}$$

then the process is positive Harris recurrent and thus admits a unique stationary measure (see [29]), where \mathcal{T}_h^{γ} is the transition kernel when we use h -rate time discretization under the piece-wise constant exploration policy γ .

The third assumption (iii) is a usual requirement for reinforcement learning algorithms.

Remark 5. The stability assumption (ii), can be replaced with so called replay buffer (as in [9]) such that if for every k , $X(k \times h) \sim \pi$, for some $\pi \in \mathcal{P}(\mathbb{X})$, the iterations (18) will converge. In particular, π will take the role of the stationary distribution under the exploration policy.

Proposition 4. Under Assumption 2, the iterations (18) converge almost surely to some $Q_h^* : \mathbb{X}_h \times \mathbb{U}_h \rightarrow \mathbb{R}$, which satisfies for every $\hat{x}, \hat{u} \in \mathbb{X}_h \times \mathbb{U}_h$

$$Q_h^*(\hat{x}, \hat{u}) = C_h^*(\hat{x}, \hat{u}) + \beta_h \sum_{\hat{x}_1 \in \mathbb{X}_h} \min_{\hat{u}_1 \in \mathbb{U}_h} Q_h^*(\hat{x}_1, \hat{u}_1) P_h^*(\hat{x}_1 | \hat{x}, \hat{u})$$

where C_h^* and P_h^* are defined in (12).

Proof. The proof can be found in Appendix B. \square

Once, Q_h^* is obtained, one can construct the policies such that

$$\gamma_h(\hat{x}) = \arg \min_{\hat{u}} Q_h^*(\hat{x}, \hat{u})$$

Using these policies, following control processes are defined

$$u_h(t) = \gamma_h(\phi_{\mathbb{X}}(X(i \times h)), \text{ for } t \in [i \times h, (i+1) \times h]). \quad (19)$$

Hence, u_h is a piece-wise constant process, which changes value at the sampling instances according to the learned map γ_h .

The the following result is a direct implication of Theorem 1, and Theorem 4:

Theorem 2. Under Assumption 1, and Assumption 2, iterations in (18) converge to some Q_h^* . For the learned control process u_h (see (19)), we have

$$W_{\beta}(x_0, u_h) - W_{\beta}^*(x_0) \leq e(h, L_{\mathbb{X}}, L_{\mathbb{U}})$$

where $e(h, L_{\mathbb{X}}, L_{\mathbb{U}})$ is defined in (17), and where u_h is the control policy learned by using the approximate Q-learning algorithm (4).

Furthermore, for small h , we have that

$$e(h, L_{\mathbb{X}}, L_{\mathbb{U}}) \leq C \left(\sqrt{h} + \frac{L_{\mathbb{X}} + L_{\mathbb{U}}}{h} + h^{\frac{1}{4}} \right)$$

for some $C < \infty$.

We now present results for the asymptotic case.

The first one states that if we first increase the quantization rate of the spaces and thus if the quantization error goes to 0, then if we take the sampling interval of the time to 0; the error bound goes to 0.

Corollary 2. *Under Assumption 1,*

$$\lim_{h \rightarrow 0} \lim_{L_{\mathbb{X}}, L_{\mathbb{U}} \rightarrow 0} (W_{\beta}(x_0, u_h) - W_{\beta}^*(x_0)) = 0.$$

We now define the quantization rate and the quantization error as a function of h and denote them by $L_{\mathbb{X}}(h), L_{\mathbb{U}}(h)$:

Corollary 3. *Under Assumption 1,*

$$\lim_{h \rightarrow 0} (W_{\beta}(x_0, u_h) - W_{\beta}^*(x_0)) = 0$$

if $L_{\mathbb{X}}(h), L_{\mathbb{U}}(h)$ go to 0 at a faster rate than h , i.e. if $\frac{L_{\mathbb{X}}(h)}{h}, \frac{L_{\mathbb{U}}(h)}{h} \rightarrow 0$ as $h \rightarrow 0$.

Remark 6. *The space discretization approach we have followed so far gives us precise error bounds and a convergence analysis under general conditions. However, we can adapt Q learning algorithms with function approximation to our setting as well. In particular, we will focus on linear approximations and discuss their convergence properties. Consider the set of Q functions that can be parametrized over the parameter θ that can be expressed as the linear span of a fixed set of M linearly independent functions $\phi_i : \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{R}$, such that the Q values can be written as*

$$Q_{\theta}(x, u) = \sum_{i=1}^M \phi_i(x, u) \theta(i),$$

then we can construct the following iterations to learn the Q values over the parametrized family:

$$\theta_{k+1} = \theta_k + \alpha_k \phi(\hat{x}_k, u_k) \Delta_k \quad (20)$$

where $\Delta_t := c(\hat{x}_t, u_t) + \beta_h \max_{u \in \mathbb{U}} Q_t(\hat{X}_{t+1}, u) - Q_t(\hat{x}_t, u_t)$, such that \hat{x}_k is the sampled diffusion process.

Then the convergence of this iterations can be shown (see [28]). In particular, one can show that the algorithm converges if

- *The sampled process converges to its stationary distribution during exploration under the exploration policy,*
- *The exploration policy is already close to the optimal policy (precise condition can be found in [28]).*

Note that the second assumption is quite restrictive. Furthermore, [28] does not provide an analysis for the error analysis of the limit Q function with respect to the optimal Q values, which in our setting would be the optimal Q values of the MDP model constructed in Section 2.1 using piece-wise constant controls for the diffusion process.

In a recent work ([9]), the assumptions are relaxed using ‘coupled Q learning’ with a so called ‘replay buffer’ assumption. If one uses the following iterations:

$$\begin{aligned} u_{t+1} &= u_t + \alpha_t(\phi(x_t, u_t)Q_{v_t}(x_t, u_t) - u_t), \\ v_{t+1} &= v_t + \beta_t\phi(x_t, u_t)\Delta_t \end{aligned}$$

where $\Delta = c(x_t, u_t) + \beta \max_{u \in \mathbb{U}} Q_{u_t}(X_{t+1}, u) - Q_{v_t}(x_t, u_t)$. It is then shown that these iterations converge if

- α_t and β_t are square summable but not summable, and $\alpha_t = o(\beta_t)$ or $\alpha_t \ll \beta_t$,
- For all t , (x_t) can be sampled from a fixed distribution, say π , or so called replay buffer.

Furthermore, for the limit Q values Q_{v^*} , we have that

$$\|Q^* - Q_{v^*}\|_\infty \leq \frac{1}{1-\beta} \|Q^* - \text{Proj}_\phi Q^*\|_\infty + \frac{1-\sigma}{\sigma} \frac{\beta\sigma}{(1-\beta)^2} \quad (21)$$

where σ is constant that depends on the set of linear basis functions. Hence, if the the optimal Q values are in the linear span of basis functions, we might get an approximation error depending on the basis functions. Note further that the replay buffer assumption replaces the stationarity assumption. However, existence of such a setup might be hard to find, i.e. one may not be able to start the process from a desired distribution.

In summary, both of these Q learning with linear approximations algorithms can be used by discretizing the time for diffusion processes. Under somehow restrictive assumptions, convergence can also be shown. Furthermore, the approximation error will be in the order of $h^{1/4}$ (piece-wise constant policy approximation error) plus the error presented in (21). However, both of these results are still not fully conclusive, as they do not analyze the performance of the learned policies but only focus on the difference between the limit Q values and the optimal Q values.

5 A Discussion on the Convergence Rate of the Learning and the Effect of Sampling and Quantization Rates on the Learning Speed

For approximation accuracy, finer sampling intervals and higher quantization rates lead to smaller error bounds. However, it is clear that finer sampling intervals and higher quantization rates also result in slower learning. In particular, higher quantization rate of the state and action spaces results in larger aggregate state and action spaces which in turn leads to dimension issues for the learning. Furthermore, finer sampling intervals increases the discount factor β_h of the approximate Q learning algorithm in (18), and higher discount rates make the iterations in (18) to converge at a slower rate by increasing the effective horizon of the problem.

First, we note that the algorithm presented in this paper relies on the convergence of the state process to its stationary distribution, since the state aggregation results in non-Markovian dynamics. Hence, the convergence speed of the algorithm depends on the convergence to the invariant measure of the process. Our motivation in this section is to study the effect of state and time discretization parameters on the speed of the Q learning algorithm. However, these discretization parameters do not affect the speed of the convergence to the stationary distribution, as this is related to the dynamics of the underlying diffusion process. Thus, for a simpler presentation, we will assume the state process starts from its stationary distribution under the exploration policy and always stays there during the exploration.

The following well known result ([10, Theorem 5]) is stated using the notation of this paper. The result provides a sample complexity bound for the near optimal Q estimates when the learning rate is linear, i.e. $\alpha_k = \frac{1}{k}$, as in this paper.

Proposition 5 ([10]). *Let Q_T be the value of the Q-learning algorithm using linear learning rate at time T . Then with probability at least $1 - \delta$, for any positive constant ψ we have $\|Q_T - Q_h^*\| \leq \epsilon$, given that*

$$T = \Omega \left((L + \psi L + 1)^{\frac{2}{1-\beta_h} \ln \frac{V_{max}}{\epsilon}} \frac{V_{max}^2 \ln(\frac{|\mathbb{X}_h| |\mathbb{U}_h| V_{max}}{\delta \epsilon \psi (1-\beta_h)})}{(\psi \epsilon (1-\beta_h))^2} \right), \quad (22)$$

where $V_{max} := \|\hat{J}_{\beta_h}\|_\infty \leq \frac{h\|c\|_\infty}{1-\beta_h}$, and L is the covering time for the algorithm, that is the smallest time for every state and action pair to be visited at least once.

The above result can be simplified for small enough h . We can write that

$$T \lesssim \left(\frac{1}{\epsilon}\right)^{\frac{\ln(L+\psi L+1)}{h}} \frac{\ln(\frac{|\mathbb{X}_h||\mathbb{U}_h|}{\delta\epsilon\psi h})}{\epsilon^2 h^2}, \quad (23)$$

where we use \lesssim as we drop some constant and logarithmic dependence.

From (23), we observe that decreasing the time discretization parameter h , increases the sample complexity exponentially. Note that, as the required sample size increases exponentially, the duration we need to observe the diffusion process in real time to get ϵ -near estimates also increases, since the exponential increase in the sample complexity dominates the decrease rate on h .

For the effect of the space discretization, we can see that the sample complexity increases in a logarithmic way with the second term, but the dominant effect is caused by the increase on the cover time L , as we have that $L \geq |\mathbb{X}_h| \times |\mathbb{U}_h|$, and depending on the sampling frequency of the samples, the cover time can be even greater. Nonetheless, the cover time increases at least linearly with the increase on the size of the aggregate state and action spaces, which in turn increases the sample complexity at a polynomial rate depending on the parameter h .

Figure 1 shows the change on the required time when for $h \in [0.7, 0.9]$, $\epsilon \in [0.1, 0.5]$ and when $|\mathbb{X}_h|$ and $|\mathbb{U}_h|$ are assumed to be order of $\frac{1}{h}$.

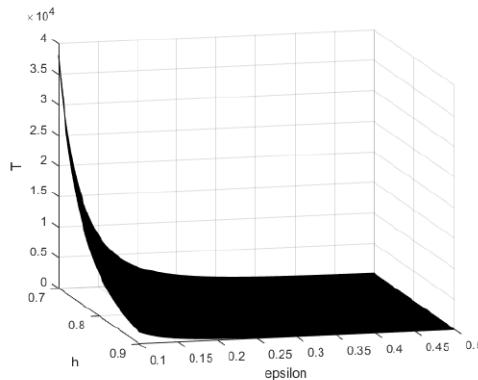


Figure 1: Sample complexity for different values of error and sampling parameters

However, this is the sample complexity to achieve ϵ -near estimates of Q_h^* which is also an estimate of the true value function (see Theorem 2 and Corollary 1). Hence, one needs to pick the sampling parameter h in a careful way considering the trade-off between the convergence speed and the approximation accuracy. Following Figure 2 provides an example for the error bounds given a given level sample points $T \sim 9500$. The graph on the left shows the difference between Q_T and Q_h^* (see (18)), clearly the error for the learning of the approximate model decays as h increases for a fixed level of sample points since the learned model becomes simpler as h increases. The graph on the right shows represents the upper bound on the difference

$$|Q_T - Q^*| \leq |Q_T - Q_h^*| + |Q_h^* - Q^*|,$$

with proper scaling. Note that the first term is the distance from the approximate Q value, whereas the second term is the approximation error. We can see that after a certain value, increasing h results on greater total error for these specific parameter intervals.

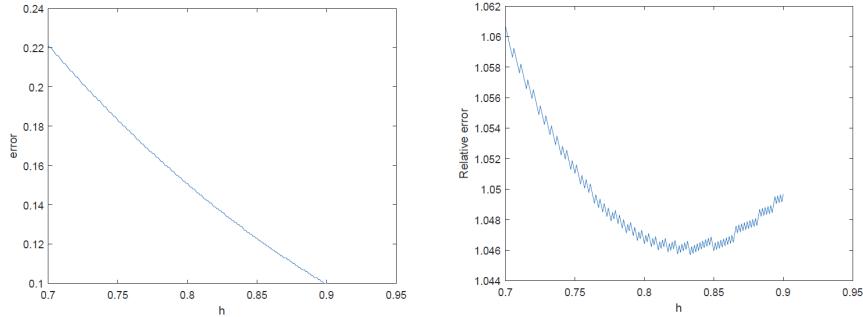


Figure 2: $|Q_T^h - Q_h^*|$ and $|Q_T^h - Q^*|$ for different values of h when $T \sim 9500$.

The sample complexity bound presented in (22) reveals that the increase is exponential in $\frac{1}{1-\beta_h}$, this rate is clearly not desired and it turns out that it can be avoided using different learning rates, α_k , rather than using linear learning rates. For example, using polynomial learning rates, $\alpha_k = \frac{1}{k^\omega}$ for some $\omega \in (1/2, 1)$, one can achieve following sample complexity for ϵ -near estimates for small enough h (see [10, Theorem 4]):

$$T \lesssim \left(\frac{L^{1+3\omega} \ln(|\mathbb{X}_h| |\mathbb{U}_h|)}{h^2 \epsilon^2} \right)^{\frac{1}{\omega}} + \left(\frac{L}{h} \ln\left(\frac{1}{\epsilon}\right) \right)^{\frac{1}{1-\omega}},$$

where the exponential increase with respect to h is eliminated.

Further improvements and variations can be achieved for the sample complexity and the convergence rate, using different learning rates, e.g. rescaled learning rates, carefully chosen constant learning rates, or with different variations of the Q learning algorithm such as speedy Q learning, or variance reduced Q learning (see e.g. [2, 44, 25]).

We note that, even though different convergence rates can be derived using different reinforcement learning algorithms, if the algorithms are constructed using the time and state discretization procedure as in this paper, the learned value function will be the value function of the approximate MDP model constructed in Section 2.2. Hence, the convergence speed can be improved with different learning rates or different Q learning variants, however,

$$\sqrt{h} + \frac{L_{\mathbb{X}} + L_{\mathbb{U}}}{h} + h^{\frac{1}{4}}$$

which is the error upper-bound for the approximation via state and time discretization and piece-wise constant policies, will not change.

6 Conclusion

We have constructed an approximate Q learning algorithm for a controlled diffusion process through discretization in time and space. We have showed that this algorithm converges under an ergodicity assumption on the state process. Furthermore, we have showed that the limit Q values satisfy the optimality equation of a finite Markov decision process, which has the same distribution as the diffusion process at the sampling points when the diffusion process is controlled with a piece-wise constant control process. Using these observations, we have derived upper bounds, as a function of the discretization parameters, for the approximation error of the learned policies compared to the performance of the optimal admissible control process.

Possible future directions, building on the analysis in this paper are as follows:

- When we discretize the state space, the aggregated state process is no longer a Markov process, hence, we use an ergodicity assumption, to guarantee the convergence of the Q learning algorithm. However, if the quantization is fine enough, one might expect the learning algorithm to stay in a set of values with sufficiently small variations even without the ergodicity assumption. Hence, a possible future problem is to relax the ergodicity assumption we consider here.

- In this paper, we have used the traditional Q learning algorithm, for the simplicity of the presentation. For faster learning rates, different variations such as the variance reduction techniques, can be considered. We note that, if one uses the same time and space quantization scheme, the learned value functions will be the same as in this paper, however, using different variations of the Q learning algorithm will change the learning speed.
- The analysis used in this paper, can be extended to partially observed systems using the results from [16, 17] under proper filter stability conditions.
- We have not considered the exploration and exploitation trade-off; one might study this relation considering the provided convergence rate in this paper.

A Proof of Proposition 1

Note that the way we constructed the controlled Markov chain in Section 2.1, implies that the state process for the controlled Markov chain, and the state process for the diffusion process have the same distribution at the sampling instances. That is

$$X_k \sim X(k \times h), \text{ for all } k \in \mathbb{Z}_+.$$

We can then write that

$$\begin{aligned} |W_\beta(x_0, u_h) - J_{\beta_h}(x_0, \gamma)| &= \left| E \left[\int_0^\infty e^{-\beta s} c(X(s), u_h(s)) ds \right] - E \left[\sum_{k=0}^\infty \beta_h^k c_h(X_k, U_k) \right] \right| \\ &= \left| E \left[\sum_{k=0}^\infty \int_{kh}^{(k+1)h} e^{-\beta s} c(X(s), u_h(s)) ds \right] - E \left[\sum_{k=0}^\infty \beta_h^k c_h(X_k, U_k) \right] \right| \\ &\leq \sum_{k=0}^\infty \int_{kh}^{(k+1)h} E \left[\left| e^{-\beta s} c(X(s), u_h(s)) - e^{-\beta hk} \frac{c_h(X_k, U_k)}{h} \right| \right] ds \end{aligned}$$

We know focus on the term inside:

$$\begin{aligned} &E \left[\left| e^{-\beta s} c(X(s), u_h(s)) - e^{-\beta hk} \frac{c_h(X_k, U_k)}{h} \right| \right] \\ &\leq E \left[\left| e^{-\beta s} c(X(s), u_h(s)) - e^{-\beta hk} c(X(s), u_h(s)) \right| \right] + E \left[\left| e^{-\beta hk} c(X(s), u_h(s)) - e^{-\beta hk} c(X_k, U_k) \right| \right] \\ &\leq e^{-\beta hk} (1 - e^{-\beta h}) \|c\|_\infty + e^{-\beta hk} K E [|X(s) - X_k|] \end{aligned}$$

for the last step, we used the fact that $u_h(s) = U_k$ as it is a piece-wise constant control process. For the second term, we have that for $s \in [kh, (k+1)h)$

$$X(s) = X_k + \int_{kh}^s b(x(r), u_h(r)) dr + \int_{kh}^s \sigma(x(r), u_h(r)) dB(r).$$

Thus, using Assumption 1 we can write

$$\begin{aligned} E [|X(s) - X_k|] &\leq \|b\|_\infty \int_{kh}^s dr + \|\sigma\|_\infty E \left[\left| \int_{kh}^s dB(r) \right| \right] \\ &\leq Bh + BE[|Z_h|] = B \left(h + \sqrt{\frac{2h}{\pi}} \right) \end{aligned}$$

where Z_h is normally distributed with mean 0 and variance h .

By combining everything we have so far, we write

$$\begin{aligned}
& |W_\beta(x_0, u_h) - J_{\beta_h}(x_0, \gamma)| \\
& \leq \sum_{k=0}^{\infty} \int_{kh}^{(k+1)h} e^{-\beta hk} (1 - e^{-\beta h}) \|c\|_\infty + e^{-\beta hk} KB \left(h + \sqrt{\frac{2h}{\pi}} \right) ds \\
& = \sum_{k=0}^{\infty} \left(e^{-\beta hk} (1 - e^{-\beta h}) \|c\|_\infty + e^{-\beta hk} KB \left(h + \sqrt{\frac{2h}{\pi}} \right) \right) h \\
& = hB + \frac{KBh}{1 - e^{-\beta h}} \left(h + \sqrt{\frac{2h}{\pi}} \right)
\end{aligned}$$

B Proof of Proposition 4

Proof. We start by writing the iterations in a more compact form by defining $C_h := c(x(i \times h), \hat{u}) \times h$, and $\hat{X}_1 := \phi_{\mathbb{X}}(X((i+1) \times h))$:

$$Q_{k+1}(\hat{x}, \hat{u}) = (1 - \alpha_k(\hat{x}, \hat{u}))Q_k(\hat{x}, \hat{u}) + \alpha_k(\hat{x}, \hat{u}) \left(C_h + \beta_h \min_{v \in \mathbb{U}_h} Q_k(\hat{X}_1, v) \right)$$

We define

$$\begin{aligned}
\Delta_k(\hat{x}, \hat{u}) &:= Q_k(\hat{x}, \hat{u}) - Q_h^*(\hat{x}, \hat{u}) \\
F_k(\hat{x}, \hat{u}) &:= C_h + \beta_h V_k(\hat{X}_1) - Q_h^*(\hat{x}, \hat{u}) \\
\hat{F}_k(\hat{x}, \hat{u}) &:= C_h^*(\hat{x}, \hat{u}) + \beta_h \sum_{\hat{x}_1} V_k(\hat{x}_1) P_h^*(\hat{x}_1 | \hat{x}, \hat{u}) - Q_h^*(\hat{x}, \hat{u}),
\end{aligned}$$

where $V_k(\hat{x}) := \min_{v \in \mathbb{U}_h} Q_k(\hat{x}, v)$.

Then, we can write the following iteration

$$\Delta_{k+1}(\hat{x}, \hat{u}) = (1 - \alpha_k(\hat{x}, \hat{u}))\Delta_k(\hat{x}, \hat{u}) + \alpha_k(\hat{x}, \hat{u})F_k(\hat{x}, \hat{u}).$$

Now, we write $\Delta_k = \delta_k + w_k$ such that

$$\begin{aligned}
\delta_{k+1}(\hat{x}, \hat{u}) &= (1 - \alpha_k(\hat{x}, \hat{u}))\delta_k(\hat{x}, \hat{u}) + \alpha_k(\hat{x}, \hat{u})\hat{F}_k(\hat{x}, \hat{u}) \\
w_{k+1}(\hat{x}, \hat{u}) &= (1 - \alpha_k(\hat{x}, \hat{u}))w_k(\hat{x}, \hat{u}) + \alpha_k(\hat{x}, \hat{u})r_k(\hat{x}, \hat{u})
\end{aligned}$$

where $r_k := F_k - \hat{F}_k = \beta_h V_k(\hat{X}_1) - \beta_h \sum_{\hat{x}_1} V_k(\hat{x}_1) P_h^*(\hat{x}_1 | \hat{x}, \hat{u}) + C_h - C_h^*(\hat{x}, \hat{u})$.

Next, we define

$$r_k^*(\hat{x}, \hat{u}) = \beta_h V^*(\hat{X}_1) - \beta_h \sum_{\hat{x}_1} V^*(\hat{x}_1) P_h^*(\hat{x}_1 | \hat{x}, \hat{u}) + C_h - C_h^*(\hat{x}, \hat{u})$$

We further separate $w_k = u_k + v_k$ such that

$$\begin{aligned} u_{k+1}(\hat{x}, \hat{u}) &= (1 - \alpha_k(\hat{x}, \hat{u}))u_k(\hat{x}, \hat{u}) + \alpha_k(\hat{x}, \hat{u})e_k(\hat{x}, \hat{u}) \\ v_{k+1}(\hat{x}, \hat{u}) &= (1 - \alpha_k(\hat{x}, \hat{u}))v_k(\hat{x}, \hat{u}) + \alpha_k(\hat{x}, \hat{u})r_k^*(\hat{x}, \hat{u}) \end{aligned}$$

where $e_k = r_k - r_k^*$.

We now show that $v_k(\hat{x}, \hat{u}) \rightarrow 0$ almost surely for all (\hat{x}, \hat{u}) . Note that, because of the way we chose the learning rates α_k , $v_k(\hat{x}, \hat{u})$ is only updated when the process hits \hat{x}, \hat{u} during the exploration. Thus, we define the following stopping times

$$\tau(n+1) = \{\min k > \tau(n) : \phi_{\mathbf{X}}(x_k) = \hat{x}, u_k = \hat{u}\}$$

where $\tau(0) = 0$. In what follows, we will focus on the $v_n(\hat{x}, \hat{u})$ process, since we assume that every (\hat{x}, \hat{u}) pair is visited infinitely often, these stopping times are bounded almost surely and hence, we can make sure that $n \rightarrow \infty$ as $k \rightarrow \infty$. Furthermore, we have that $\alpha_n(\hat{x}, \hat{u}) = \frac{1}{n}$, as n is the number of times the we have hit (\hat{x}, \hat{u}) pair.

When $\alpha_n(\hat{x}, \hat{u}) = \frac{1}{n}$ for every (\hat{x}, \hat{u}) pair, the problem reduces to

$$v_{n+1}(\hat{x}, \hat{u}) = \frac{1}{n} \sum_{n'=0}^{n-1} r_{n'}^*(\hat{x}, \hat{u}).$$

Above, even though, we do not write the time dependence on \hat{X}_1 , distribution of the \hat{X}_1 is different every time we make the update, since the marginal distribution of x_t is different. However, using the fact that the random variables (\hat{X}_1, X_1) form a controlled Markov chain since they are sampled at the stopping times $\tau(k)$, and using the invariant measure of the original state process, we can write

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{n'=0}^{n-1} V^*(\hat{X}_1) &= \int_B \sum_j V^*(\hat{x}_j) \mathcal{T}_h(B_j | x, \hat{u}) \hat{\pi}_{\hat{x}}^*(dx) \\ \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{n'=0}^{n-1} C_h &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{n'=0}^{n-1} c(x(\tau(n') \times h), \hat{u}) \times h = \int_B c_h(x, u) \hat{\pi}_{\hat{x}}^*(dx). \end{aligned}$$

where B is the quantization bin \hat{x} belongs to. Hence, we have proved that $v_k(\hat{x}, \hat{u}) \rightarrow 0$, noting

$$\begin{aligned} C_h^*(\hat{x}_i, u) &= \int_{B_i} c_h(x, u) \hat{\pi}_{\hat{x}_i}^*(dx) \\ P_h^*(\hat{x}_j | \hat{x}_i, u) &= \int_{B_i} \mathcal{T}_h(B_j | x, u) \hat{\pi}_{\hat{x}_i}^*(dx). \end{aligned}$$

Now, we go back to the iterations:

$$\begin{aligned}\delta_{k+1}(\hat{x}, \hat{u}) &= (1 - \alpha_k(\hat{x}, \hat{u}))\delta_k(\hat{x}, \hat{u}) + \alpha_k(\hat{x}, \hat{u})\hat{F}_k(\hat{x}, \hat{u}) \\ u_{k+1}(\hat{x}, \hat{u}) &= (1 - \alpha_k(\hat{x}, \hat{u}))u_k(\hat{x}, \hat{u}) + \alpha_k(\hat{x}, \hat{u})e_k(\hat{x}, \hat{u}) \\ v_{k+1}(\hat{x}, \hat{u}) &= (1 - \alpha_k(\hat{x}, \hat{u}))v_k(\hat{x}, \hat{u}) + \alpha_k(\hat{x}, \hat{u})r_k^*(\hat{x}, \hat{u}).\end{aligned}$$

Note that, we want to show $\Delta_k = \delta_k + u_k + v_k \rightarrow 0$ almost surely and we have that $v_k(\hat{x}, \hat{u}) \rightarrow 0$ almost surely for all (\hat{x}, \hat{u}) . The following analysis holds for any path that belongs to the probability one event in which $v_k(I, u) \rightarrow 0$. For any such path and for any given $\epsilon > 0$, we can find an $N < \infty$ such that $\|v_k\|_\infty < \epsilon$ for all $k > N$ as (I, u) takes values from a finite set.

We now focus on the term $\delta_k + u_k$ for $k > N$:

$$(\delta_{k+1} + u_{k+1})(\hat{x}, \hat{u}) = (1 - \alpha_k(\hat{x}, \hat{u}))(\delta_k + u_k)(\hat{x}, \hat{u}) + \alpha_k(\hat{x}, \hat{u})(\hat{F}_k + e_k)(\hat{x}, \hat{u}). \quad (24)$$

Observe that for $k > N$,

$$\begin{aligned}(\hat{F}_k + e_k)(\hat{x}, \hat{u}) &= (F_k - r_k^*)(\hat{x}, \hat{u}) = \beta_h V_k(\hat{X}_1) - \beta_h V^*(\hat{X}_1) \\ &\leq \beta_h \max_{\hat{x}, \hat{u}} |Q_k(\hat{x}, \hat{u}) - Q^*(\hat{x}, \hat{u})| = \beta_h \|\Delta_k\|_\infty \\ &\leq \beta_h \|\delta_k + u_k\|_\infty + \beta_h \epsilon\end{aligned}$$

where the last step follows from the fact that $v_k \rightarrow 0$ almost surely. By choosing $C < \infty$ such that $\hat{\beta} := \beta_h(C+1)/C < 1$, for $\|\delta_k + u_k\|_\infty > C\epsilon$, we can write that

$$\beta_h \|\delta_k + u_k + \epsilon\|_\infty \leq \hat{\beta} \|\delta_k + u_k\|_\infty.$$

Now we rewrite (24)

$$\begin{aligned}(\delta_{k+1} + u_{k+1})(\hat{x}, \hat{u}) &= (1 - \alpha_k(\hat{x}, \hat{u}))(\delta_k + u_k)(\hat{x}, \hat{u}) + \alpha_k(\hat{x}, \hat{u})(\hat{F}_k + e_k)(\hat{x}, \hat{u}) \\ &\leq (1 - \alpha_k(\hat{x}, \hat{u}))(\delta_k + u_k)(\hat{x}, \hat{u}) + \alpha_k(\hat{x}, \hat{u})\hat{\beta} \|\delta_k + u_k\|_\infty.\end{aligned}$$

By [13, Lemma 3], $(\delta_{k+1} + u_{k+1})(\hat{x}, \hat{u})$ tends to 0 for $\|\delta_k + u_k\|_\infty > C\epsilon$. This shows that the condition $\|\delta_k + u_k\|_\infty > C\epsilon$ cannot be sustained indefinitely. Next, we show that once the process hits below $C\epsilon$ it always stays there. Suppose $\|\delta_k + u_k\|_\infty < C\epsilon$,

$$\begin{aligned}(\delta_{k+1} + u_{k+1})(\hat{x}, \hat{u}) &\leq (1 - \alpha_k(\hat{x}, \hat{u}))(\delta_k + u_k)(\hat{x}, \hat{u}) + \alpha_k(\hat{x}, \hat{u})\beta_h (\|\delta_k + u_k\|_\infty + \epsilon) \\ &\leq (1 - \alpha_k(\hat{x}, \hat{u}))C\epsilon + \alpha_k(\hat{x}, \hat{u})\beta_h(C\epsilon + \epsilon) \\ &= (1 - \alpha_k(\hat{x}, \hat{u}))C\epsilon + \alpha_k(\hat{x}, \hat{u})\beta_h(C+1)\epsilon \\ &\leq (1 - \alpha_k(\hat{x}, \hat{u}))C\epsilon + \alpha_k(\hat{x}, \hat{u})C\epsilon, \quad (\beta_h(C+1) \leq C) \\ &= C\epsilon.\end{aligned}$$

Then, we can write $\|\delta_{k+1} + u_{k+1}\|_\infty < C\epsilon$.

Thus, taking $\epsilon \rightarrow 0$, we can conclude that $\Delta_k = \delta_k + u_k + v_k \rightarrow 0$ almost surely.

Therefore, the process Q_k , determined by the algorithm converges almost surely to Q_h^* . \square

C Proof of Lemma 1

We are interested in the distance between the distributions of the following random variables

$$\begin{aligned} x(h) &= x(0) + \int_0^h b(x(s), u) ds + \int_0^h \sigma(x(s), u) dB(s) \\ y(h) &= y(0) + \int_0^h b(y(s), u) ds + \int_0^h \sigma(y(s), u) dB(s). \end{aligned}$$

We define $z(h) := x(h) - y(s)$, whose dynamics are given by

$$z(h) = z(0) + \int_0^h \bar{b}(x(s), y(s), u) ds + \int_0^h \bar{\sigma}(x(s), y(s), u) dB(s)$$

where

$$\begin{aligned} \bar{b}(x(s), y(s), u) &:= b(x(s), u) - b(y(s), u) \\ \bar{\sigma}(x(s), y(s), u) &:= \sigma(x(s), u) - \sigma(y(s), u). \end{aligned}$$

Using the Ito formula

$$\begin{aligned} E[Z(h)^2] &\leq Z_0^2 + E \left[\int_0^h 2Z_t \bar{b}(x(s), y(s), u) ds \right] + E \left[\int_0^h \bar{\sigma}^2(x(s), y(s), u) ds \right] \\ &\leq Z_0^2 + 2K \int_0^h E[Z_s^2] ds + K^2 \int_0^h E[Z_s^2] ds. \end{aligned}$$

We can then use Gronwall's inequality to write

$$E[Z(h)^2] \leq Z_0^2 e^{(2K+K^2)h}.$$

Using the Holder's inequality, we can further write

$$E[|X(h) - Y(h)|] \leq |x(0) - y(0)| e^{(K+\frac{K^2}{2})h},$$

which concludes the proof.

D Proof of Lemma 2

We first define the following random variables

$$\begin{aligned} x(h) &= x(0) + \int_0^h b(x(s), u) ds + \int_0^h \sigma(x(s), u) dB(s) \\ y(h) &= x(0) + \int_0^h b(y(s), \hat{u}) ds + \int_0^h \sigma(y(s), \hat{u}) dB(s). \end{aligned}$$

By defining $z(h) := x(h) - y(s)$, whose dynamics are given by

$$z(h) = \int_0^h \bar{b}(x(s), y(s), u, \hat{u}) ds + \int_0^h \bar{\sigma}(x(s), y(s), u, \hat{u}) dB(s)$$

where

$$\begin{aligned} \bar{b}(x(s), y(s), u, \hat{u}) &:= b(x(s), u) - b(y(s), \hat{u}) \\ \bar{\sigma}(x(s), y(s), u, \hat{u}) &:= \sigma(x(s), u) - \sigma(y(s), \hat{u}). \end{aligned}$$

Using the Holder's inequality, the boundedness, and the continuity properties of b, σ , under the assumption that $h < 1$, we can write

$$\begin{aligned} E[Z(h)^2] &\leq E \left[\int_0^h (KZ(s) + K|u - \hat{u}|)^2 ds \right] + E \left[\left(\int_0^h (KZ(s) + K|u - \hat{u}|) dB(s) \right)^2 \right] \\ &= 2E \left[\int_0^h (KZ(s) + K|u - \hat{u}|)^2 ds \right] \leq 4K^2 \int_0^h E[Z(s)^2] ds + 4K^2 h |u - \hat{u}|^2 \end{aligned}$$

under the assumption that $h < 1$, and using the Gronwall inequality

$$E[Z(h)^2] \leq 4K^2 |u - \hat{u}|^2 e^{4K^2 h}.$$

Using the Holder's inequality, we can conclude that

$$E[|X(h) - Y(h)|] \leq 2K e^{2K^2 h} |u - \hat{u}|.$$

References

- [1] Murad Abu-Khalaf and Frank L Lewis. Nearly optimal control laws for nonlinear systems with saturating actuators using a neural network hjb approach. *Automatica*, 41(5):779–791, 2005.
- [2] M. G. Azar, R. Munos, M. Ghavamzadeh, and H. J. Kappen. Speedy q-learning. In *NIPS*, 2011.

- [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] G. Barles and E.R. Jakobsen. Error bounds for monotone approximation schemes for Hamilton-Jacobi-Bellman equations. *SIAM J. Numer. Anal.*, 43(2):540–558.
- [5] G. Barles and E.R. Jakobsen. On the convergence rate of approximation schemes for Hamilton-Jacobi-Bellman equations. *Math. Model. Numer. Anal.*, 36:33–54.
- [6] Philipp Batz, Andreas Rutter, and Manfred Opper. Approximate bayes learning of stochastic differential equations. *Physical Review E*, 98(2):022109, 2018.
- [7] Erhan Bayraktar and Ali Devran Kara. An approximate reinforcement learning algorithm for controlled diffusion processes. *arXiv preprint arXiv:2203.07499*, 2022.
- [8] Tao Bian and Zhong-Ping Jiang. Value iteration and adaptive dynamic programming for data-driven adaptive optimal control design. *Automatica*, 71:348–360, 2016.
- [9] Diogo Carvalho, Francisco S Melo, and Pedro Santos. A new convergent variant of q-learning with linear function approximation. *Advances in Neural Information Processing Systems*, 33:19412–19421, 2020.
- [10] E. Even-Dar, Y. Mansour, and P. Bartlett. Learning rates for q-learning. *Journal of machine learning Research*, 5(1), 2003.
- [11] C. Gaskett and A. Zelinsky D. Wettergreen. Q-learning in continuous state and action spaces. In *Australasian joint conference on artificial intelligence*, pages 417–428. Springer, 1999.
- [12] Maximilien Germain, Huyêñ Pham, and Xavier Warin. Neural networks-based algorithms for stochastic control and pdes in finance. *arXiv preprint arXiv:2101.08068*, 2021.
- [13] T. Jaakkola, M. I. Jordan, and S. P. Singh. On the convergence of stochastic iterative dynamic programming algorithms. *Neural computation*, 6(6):1185–1201, 1994.

- [14] E. R. Jakobsen, A. Picarelli, and C. Reisinger. Improved order 1/4 convergence for piecewise constant policy approximation of stochastic control problems. *Electronic Communications in Probability*, 24:1–10, 2019.
- [15] A. D. Kara, N. Saldi, and S. Yüksel. Q-learning for mdps with general spaces: Convergence and near optimality via quantization under weak continuity. *arXiv preprint arXiv:2111.06781*, 2021.
- [16] A. D. Kara and S. Yuksel. Convergence of finite memory q-learning for pomdps and near optimality of learned policies under filter stability. *arXiv preprint arXiv:2103.12158*, 2021.
- [17] A. D. Kara and S. Yuksel. Near optimality of finite memory feedback policies in partially observed markov decision processes. *Journal of Machine Learning Research*, 23(1):1–46, 2022.
- [18] George P Kontoudis and Kyriakos G Vamvoudakis. Kinodynamic motion planning with continuous-time q-learning: An online, model-free, and safe navigation framework. *IEEE transactions on neural networks and learning systems*, 30(12):3803–3817, 2019.
- [19] N. V. Krylov. On the rate of convergence of finite-difference approximations for Bellman’s equations. *St. Petersburg Math. J.*, 9:639–650.
- [20] N. V. Krylov. On the rate of convergence of finite-difference approximations for Bellman’s equations with variable coefficients. *Probab. Theory Relat. Fields*, 117:1–16.
- [21] N. V. Krylov. Approximating value functions for controlled degenerate diffusion processes by using piece-wise constant policies. *Electronic Journal of Probability*, 4:1–19, 1999.
- [22] H.J. Kushner and P.G. Dupuis. *Numerical methods for stochastic control problems in continuous time*, volume 24. Springer Science & Business Media, 2001.
- [23] Jaeyoung Lee and Richard S Sutton. Policy iterations for reinforcement learning problems in continuous time and space—fundamental theory and methods. *Automatica*, 126:109421, 2021.
- [24] William Lefebvre, Grégoire Loeper, and Huyêñ Pham. Differential learning methods for solving fully nonlinear pdes. *arXiv preprint arXiv:2205.09815*, 2022.

- [25] G. Li, Y. Wei, Y. Chi, Y. Gu, and Y. Chen. Sample complexity of asynchronous q-learning: Sharper analysis and variance reduction. *arXiv preprint arXiv:2006.03041*, 2020.
- [26] Shujian Liao, Terry Lyons, Weixin Yang, and Hao Ni. Learning stochastic differential equations using rnn with log signature features. *arXiv preprint arXiv:1908.08286*, 2019.
- [27] Michael Lutter, Boris Belousov, Kim Listmann, Debora Clever, and Jan Peters. Hjb optimal feedback control with deep differential value functions and action constraints. In *Conference on Robot Learning*, pages 640–650. PMLR, 2020.
- [28] F. C. Melo, S. P. Meyn, and I. M. Ribeiro. An analysis of reinforcement learning with function approximation. In *Proceedings of the 25th international conference on Machine learning*, pages 664–671, 2008.
- [29] S. P. Meyn and R. Tweedie. *Markov Chains and Stochastic Stability*. Springer-Verlag, London, 1993.
- [30] R. Munos and P. Bourgine. Reinforcement learning for continuous stochastic control problems. *Advances in neural information processing systems*, 10, 1997.
- [31] Rémi Munos. Policy gradient in continuous time. *Journal of Machine Learning Research*, 7:771–791, 2006.
- [32] Muthukumar Palanisamy, Hamidreza Modares, Frank L Lewis, and Muhammad Aurangzeb. Continuous-time q-learning for infinite-horizon discounted cost linear quadratic regulator problems. *IEEE transactions on cybernetics*, 45(2):165–176, 2014.
- [33] José Pereira, Morteza Ibrahimi, and Andrea Montanari. Learning networks of stochastic differential equations. *Advances in Neural Information Processing Systems*, 23, 2010.
- [34] Karthikeyan Rajagopal, Sivasubramanya Nadar Balakrishnan, and Jerome R Busemeyer. Neural network-based solutions for stochastic optimal control using path integrals. *IEEE transactions on neural networks and learning systems*, 28(3):534–545, 2016.
- [35] D. Shah and Q. Xie. Q-learning with nearest neighbors. *arXiv preprint arXiv:1802.03900*, 2018.

- [36] S. P. Singh, T. Jaakkola, and M. I. Jordan. Reinforcement learning with soft state aggregation. *Advances in neural information processing systems*, pages 361–368, 1995.
- [37] C. Szepesvári. Algorithms for reinforcement learning. volume 4, pages 1–103, 2010.
- [38] C. Szepesvári and William D. Smart. Interpolation-based q-learning.. 2004.
- [39] C. Tallec, L. Blier, and Y. Ollivier. Making deep q-learning methods robust to time discretization. In *International Conference on Machine Learning*, pages 6096–6104. PMLR, 2019.
- [40] Yuval Tassa and Tom Erez. Least squares solutions of the hjb equation with neural network value-function approximators. *IEEE transactions on neural networks*, 18(4):1031–1041, 2007.
- [41] J. N. Tsitsiklis. Asynchronous stochastic approximation and q-learning. *Machine Learning*, 16:185–202, 1994.
- [42] J. N. Tsitsiklis and B. Van Roy. An analysis of temporal-difference learning with function approximation. *IEEE transactions on automatic control*, 42(5):674–690, 1997.
- [43] Kyriakos G Vamvoudakis. Q-learning for continuous-time linear systems: A model-free infinite horizon optimal control approach. *Systems & Control Letters*, 100:14–20, 2017.
- [44] M. Wainwright. Stochastic approximation with cone-contractive operators: Sharper l_∞ -bounds for q -learning. *arXiv preprint arXiv:1905.06265*, 2019.
- [45] H. Wang, T. Zariphopoulou, and X. Y. Zhou. Reinforcement learning in continuous time and space: A stochastic control approach. *J. Mach. Learn. Res.*, 21(198):1–34, 2020.
- [46] H. Wang and X. Y. Zhou. Continuous-time mean–variance portfolio selection: A reinforcement learning framework. *Mathematical Finance*, 30(4):1273–1308, 2020.
- [47] C. J. C. H. Watkins and P. Dayan. Q-learning. *Machine Learning*, 8:279–292, 1992.

- [48] Yongliang Yang, Donald Wunsch, and Yixin Yin. Hamiltonian-driven adaptive dynamic programming for continuous nonlinear dynamical systems. *IEEE transactions on neural networks and learning systems*, 28(8):1929–1940, 2017.
- [49] Cagatay Yildiz, Markus Heinonen, Jukka Intosalmi, Henrik Mannerstrom, and Harri Lähdesmäki. Learning stochastic differential equations with gaussian processes without gradient matching. In *2018 IEEE 28th International Workshop on Machine Learning for Signal Processing (MLSP)*, pages 1–6. IEEE, 2018.
- [50] Cagatay Yildiz, Markus Heinonen, and Harri Lähdesmäki. Continuous-time model-based reinforcement learning. In *International Conference on Machine Learning*, pages 12009–12018. PMLR, 2021.