



HAL
open science

On a Variance Reduction Correction of the Temporal Difference for Policy Evaluation in the Stochastic Continuous Setting

Ziad Kobeissi, Francis Bach

► **To cite this version:**

Ziad Kobeissi, Francis Bach. On a Variance Reduction Correction of the Temporal Difference for Policy Evaluation in the Stochastic Continuous Setting. 2022. hal-03574645v1

HAL Id: hal-03574645

<https://inria.hal.science/hal-03574645v1>

Preprint submitted on 15 Feb 2022 (v1), last revised 5 Jun 2023 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On a Variance Reduction Correction of the Temporal Difference for Policy Evaluation in the Stochastic Continuous Setting

Ziad Kobeissi
Institut Louis Bachelier
Inria Paris
ziad.kobeissi@inria.fr

Francis Bach
Inria & École Normale Supérieure
PSL Research University
francis.bach@inria.fr

February 15, 2022

Abstract

This paper deals with solving continuous time, state and action optimization problems in stochastic settings, using reinforcement learning algorithms, and considers the policy evaluation process. We prove that standard learning algorithms based on the discretized temporal difference are doomed to fail when the time discretization tends to zero, because of the stochastic part. We propose a variance-reduction correction of the temporal difference, leading to new learning algorithms that are stable with respect to vanishing time steps. This allows us to give theoretical guarantees of convergence of our algorithms to the solutions of continuous stochastic optimization problems.

1 Introduction

Most of the literature in reinforcement learning (RL) is dedicated to the discrete setting in time. Given a state set Ω and an admissible action set \mathcal{U} , it aims at finding an action function $u : \Omega \rightarrow \mathcal{U}$ maximizing the cumulative reward $R(u)$ defined by

$$(1.1) \quad R(u) = \mathbb{E} \left[\sum_{i=0}^{\infty} \gamma^i r(x_i, u(x_i)) \right], \quad \text{with } x_{i+1} = b(x_i, u(x_i)),$$

where $\gamma \in (0, 1)$ is the discount factor, $r : \Omega \times \mathcal{U} \rightarrow \mathbb{R}$ is the instant reward function, $x_i \in \Omega$ is the state at time i , and $b : \Omega \times \mathcal{U} \rightarrow \Omega$ represents the dynamics. Depending on the set-up, the functions r and b might be fully or partially known, deterministic or stochastic.

A standard starting point for solving the latter discrete optimization problem consists in introducing the value function $V^u : \Omega \rightarrow \mathbb{R}$, for a specific action function u , defined by

$$V^u(x) = \mathbb{E} \left[\sum_{i=0}^{\infty} \gamma^i r(x_i, u(x_i)) \middle| x_0 = x \right].$$

Reinforcement learning is often used to approximate solutions of continuous problems, using time discretization. The discretization error has already been studied in the deterministic set-up (Dayan and Singh, 1996; Doya, 2000; Lutter et al., 2021). However, the generalization to the stochastic case is in general missing. This is what we propose to do here. Let us introduce a stochastic continuous counterpart to the discrete formulation. The state is now denoted by $(X_t)_{t \in [0, \infty)}$, a random process in Ω which satisfies the following stochastic differential equation (SDE),

$$(1.2) \quad dX_t = b(X_t, u(X_t))dt + \sigma(X_t, u(X_t))dW_t,$$

where W is a d -dimensional Brownian motion and σ is a matrix-valued function. The continuous value function is defined by (with $\rho > 0$ the continuous discount factor):

$$(1.3) \quad V^u(x) = \mathbb{E} \left[\int_0^\infty e^{-\rho t} r(X_t, u(X_t)) dt \mid X_0 = x \right].$$

Iterative algorithms based on value functions or similar concepts are known as generalized policy iteration (GPI) methods (Sutton and Barto, 2018). They consist in alternating two interactive updates: the policy evaluation and the policy improvement. In the present work, we focus on the policy evaluation process. It consists in computing an approximation of the value function V^u defined in (1.3), assuming that we can access $u(x)$ for any given $x \in \Omega$.

Concerning the policy improvement process, which will not be treated in the present work, it consists in changing the policy in order to increase the cumulative reward. For instance, using dynamic programming given an approximate value function V , an admissible policy update consists in computing u defined by the following closed form:

$$u(x) = \operatorname{argmax}_{u' \in \mathcal{U}} \left\{ r(x, u') + b(x, u') \cdot \nabla_x V(x) + \frac{1}{2} \operatorname{tr} \left((\sigma \sigma^\top)(x, u') D_{xx}^2 V(x) \right) \right\}.$$

More classical results on stochastic control theory are given in Section 2.1. Alternatively, if the model is not known, it is not possible to compute the latter control function, and one may use for instance an actor-critic method (Konda and Tsitsiklis, 2000).

The simplest idea to approximate the solution of such an infinite dimensional optimization problem is to discretize the continuous dynamics and rewards with respect to time with a sufficiently small time step $\Delta t > 0$. The simplest discretization scheme is the Euler-Maruyama scheme (Kloeden and Platen, 1992), that we adopt here for simplicity (however the arguments of this work extend straightforwardly to higher-order discretization schemes),

$$(1.4) \quad X_{t_{i+1}} = S_{\Delta t}(X_{t_i}, \xi_i) := X_{t_i} + \Delta t b(X_{t_i}, u(X_{t_i})) + \sqrt{\Delta t} \sigma(X_{t_i}, u(X_{t_i})) \xi_i,$$

where $S_{\Delta t}$ is the step operator, $t_i = i\Delta t$ for $i \in \mathbb{N}$, and $(\xi_i)_{i \geq 0}$ are independent normally distributed random variables with mean zero and covariance identity. For readers who are not familiar with Brownian motion, when Δt tends to zero, the fact that during a single step the drift $\Delta t b$ in (1.4) becomes negligible with respect to the magnitude of the noise $\sqrt{\Delta t} \sigma \xi$, may seem surprising. However, this is necessary for the cumulative noise (defined by $\sum_{i \text{ s.t. } t \leq t_i < t'} \sqrt{\Delta t} \sigma(X_{t_i}, u(X_{t_i})) \xi_i$ for $0 < t < t'$) not to vanish at convergence because of the law of large numbers. Then, a Brownian motion is simply a continuous limit of the latter discrete noise processes. This remains if $(\xi_i)_{i \geq 0}$ are not Gaussian random variables, as long as they are independent centered random variables with bounded variances, as a consequence of Donsker's theorem (Donsker, 1951).

1.1 Motivating examples

The specific form of the noise being motivated in the previous paragraph, we now motivate it by presenting a list of three classes of models for applications:

- C1** stochastic models of the form of (1.2) with an intrinsic noise,
- C2** deterministic models with linear dynamics with respect to the controls, and an artificial noise added for exploration, regularization or for smoothing the control (Le Lidec et al., 2021).
- C3** stochastic models of the form of (1.2), with linear dynamics with respect to the controls, and an artificial noise.

For the first class of models **C1**, the noise is part of the model and cannot be tuned. The last two classes seem more interesting in the framework of RL, and more specifically in the theoretical study of the exploration/exploitation trade-off. In class **C2**, the dynamics has the following form,

$$(1.5) \quad \frac{d}{dt}x_t = A(x_t)u(x_t) + B(x_t),$$

where A and B are respectively matrix-valued and vector-valued functions. Then, in order to encourage exploration, instead of choosing a deterministic control function (being greedy with respect to some criterion), one generally adds noises in the choice of u . Gaussian noises are often considered in discrete dynamics because of their simplicity to sample, or because they are the minimizers of some entropy-relaxations of the optimization problems (see Wang et al. (2020) for instance). Therefore, at least at the discrete level, it is natural to change u into its noisy counterpart $u + \sigma(x, u)\xi_i/\sqrt{\Delta t}$. This leads to the following dynamics,

$$X_{t_{i+1}} = X_{t_i} + \Delta t (A(X_{t_i})u(X_{t_i}) + B(X_{t_i})) + \sqrt{\Delta t}A(X_{t_i})\sigma(X_{t_i}, u(X_{t_i}))\xi_i,$$

which admits a similar form as (1.4) with $A\sigma$ replacing σ . This time, the noise is tunable and a particular interesting regime consists in letting σ tends to zero. The class **C3** consists in a mix between the two other classes, we also get a similar discrete dynamics as (1.4) with σ replaced by $A\sigma_{\text{art}} + \sigma_{\text{int}}$, where σ_{art} and σ_{int} are the artificial and intrinsic noises respectively. The noise is tunable in some measure but the regime $\sigma \rightarrow 0$ is in general prohibited.

1.2 Temporal differences and variance-reduction corrections

Coming back to the general case of a discretized optimization problem with dynamics given in (1.4), we are facing a problem in a similar form as (1.1), where b in the discrete problem is replaced by $\Delta t b + \sigma\sqrt{\Delta t}\xi$, r is replaced by $\Delta t r$, and $\gamma = e^{-\rho\Delta t}$. Coupling this with function approximation, we introduce the parametrized learnt value function $v(\cdot, \theta)$, for $\theta \in \Theta$ and Θ the set of admissible parameters. This leads to define the rescaled temporal difference as

$$(1.6) \quad \tilde{\delta}_{\Delta t}(X_{t_i}, \xi_i, \theta) = (\Delta t)^{-1}(v(X_{t_i}, \theta) - \gamma v(S_{\Delta t}(X_{t_i}, \xi_i), \theta) - r(X_{t_i}, u(X_{t_i}))\Delta t),$$

where the normalization constant $(\Delta t)^{-1}$ is chosen in the latter definition because $\mathbb{E}[\tilde{\delta}_{\Delta t}]$ is convergent to its continuous counterpart when Δt tends to 0. Choosing a different order of magnitude in the normalization constant would be pointless as it leads to a convergence in average to zero or infinity.

At least heuristically, we get the following Taylor expansion of $\tilde{\delta}_{\Delta t}$ when $\Delta t \rightarrow 0$,

$$\begin{aligned} \tilde{\delta}_{\Delta t} &= (\Delta t)^{-1} \left[v - (1 - \rho\Delta t) \left(v + \nabla_x v^\top (\Delta t b + \sqrt{\Delta t}\sigma\xi) + \Delta t \xi^\top \sigma^\top (D_x^2 v)\sigma\xi \right) - \Delta t r \right] + o(1) \\ &= \rho v - \nabla_x v^\top b - (\Delta t)^{-\frac{1}{2}} \nabla_x v^\top \sigma\xi - \xi^\top \sigma^\top (D_x^2 v)\sigma\xi - r + o(1), \end{aligned}$$

where we omitted the arguments in the functions $\tilde{\delta}_{\Delta t}$, v , b , σ and r , to simplify the notations. One may notice that the third term in the right-hand side of the latter equality has zero mean when ξ is a centered random variable, but its variance is of order Δt^{-1} . We thus obtain that any RL algorithm based on such a temporal difference is doomed to fail when $\Delta t \rightarrow 0$ because of the latter variance tending to infinity. Our first contribution consists in proving this claim in the particular cases of the residual gradient method (Baird, 1995), and TD(0) (Sutton, 1988). We are confident that similar conclusions hold for other algorithms such as exponential eligibility trace TD(λ) (Sutton, 1988), Q-learning, actor-critic, and other algorithms based on the temporal difference.

Our second contribution is to propose a correction to the temporal difference; it consists in adding an additional term to the rescaled temporal difference to overcome the difficulties above.

More precisely, we define the rescaled *corrected* temporal difference as

$$(1.7) \quad \begin{aligned} \delta_{\Delta t} &= (\Delta t)^{-1} (v(X_{t_i}, \theta) - \gamma v(S_{\Delta t}(X_{t_i}, \xi_i), \theta) - r(X_{t_i}, u(X_{t_i}))\Delta t + Z_{\Delta t}(X_{t_i}, \xi_i, \theta)), \\ \text{where } Z_{\Delta t}(X_{t_i}, \xi_i, \theta) &= \sqrt{\Delta t} \nabla_x v(X_{t_i}, \theta)^\top \sigma(X_{t_i}, u(X_{t_i})) \xi_i. \end{aligned}$$

This allows us to propose two adaptations of the residual gradient algorithm (and three more in Section 4.1), and one of TD(0), and prove that they all converge to a desired continuous solution.

We refer to the additional term $Z_{\Delta t}$ as the variance-reduction term. This terminology is due to the facts that the conditional expectation with respect to X_{t_i} of $Z_{\Delta t}$ is equal to zero, and that the variance of $\tilde{\delta}_{\Delta t}$ is of order $(\Delta t)^{-1}$ while the variance of $\delta_{\Delta t}$ is uniformly bounded. The use of the letter Z for the notation comes from its continuous counterpart appearing in the backward stochastic differential equation (BSDE) point of view of the optimal control problem (see Pham et al. (2021) for instance). Here, to compute Z , we assume that we can access $\sqrt{\Delta t}\sigma\xi$; this makes the algorithms presented in this work being *model-based*. Alternatively, we may assume that only b can be accessed, and we compute Z by the following characterization,

$$Z_{\Delta t}(X_{t_i}, \xi_i, \theta) = \tilde{Z}_{\Delta t}(X_{t_i}, X_{t_{i+1}}, \theta) := (X_{t_{i+1}} - X_{t_i} - b(X_{t_i}, u(X_{t_i}))\Delta t) \cdot \nabla_x v(X_{t_i}, \theta).$$

One may notice that the noise does not have to be normally distributed in the latter formulation. Therefore, the arguments in this work can easily be extended to other noise distributions.

It is also possible to extend our algorithms to be model-free by learning $Z_{\Delta t}/\sqrt{\Delta t}$ instead of computing it. Examples of such learning algorithms may be adapted from the survey of Pham et al. (2021) on learning approximations of solutions of backward stochastic differential equations.

1.3 Related works

Continuous-time reinforcement learning started with the seminal work of Baird III (1993), who proposed a continuous-time counterpart to Q -learning; it was later extended by Tallec et al. (2019). From a different perspective, Bradtke and Duff (1994) extended classical RL algorithms to continuous-time discrete-state Markov decision processes. Then, using deterministic dynamics given by ordinary differential equations (ODE), and based on the Hamilton-Jacobi-Bellman (HJB) equation, Doya (2000) derived algorithms for both policy evaluation and policy improvement. This deterministic approach of continuous-time RL has then be explored recently by Lutter et al. (2021); Yildiz et al. (2021). In order to balance between exploration and exploitation, Wang et al. (2020) added an entropy-regularization term to a similar continuous optimization problem, they concluded that Gaussian controls are optimal for their relaxed problems, leading to a similar SDE system as the one studied in the present work.

One novelty of the present work is that we prove convergence results with explicit convergence rates of stochastic iterative methods to approximating solutions of the optimization problem in the stochastic continuous set-up. Those convergence results may be compared to standard results on RL algorithm, see Bellman (1966); Kirk (1970) for TD(0), and Baird (1995) for the residual gradient algorithm. The techniques in the proof (especially concerning the fast-convergence results in Section 3.3) are adapted from the literature on stochastic gradient descent methods (Polyak and Juditsky, 1992; Bach and Moulines, 2013) to the non symmetric setting.

1.4 Summary of contributions

Our first contribution is to give theoretical evidences that learning methods based on the temporal difference, are not adapted to high-frequency optimization. More precisely, those methods are doomed to fail when applied to the discretization of a continuous stochastic problem, when we let the discretization step tends to zero. This claim is made clear by the first equality in Lemma 2.3

in Section 2.2, where it is shown that the variance of the standard rescaled temporal difference tend to infinity, making learning impossible.

Our second contribution consists in proposing a correction to the temporal difference, based on the Taylor expansion at a neighborhood of the continuous problem when the time step tends to zero. Namely, the variance of the corrected rescaled temporal difference stays bounded at the limit, see the second equality in Lemma 2.3.

Our third contribution is to prove the convergence of adaptations of well-known algorithms, TD(0) and the residual gradient method, which do not hold without the variance-reduction correction. First, using usual decreasing learning steps and relaxation parameters, we prove standard convergence rates for our algorithms, see Theorem 3.2 in Section 3.2. Second, without relaxation and with constant learning step, we prove a convergence result using a Polyak-Juditsky averaging method, see Theorem 3.3 in Section 3.3. Our rate of convergence, of order $1/k$, is analogous to the optimal rate of convergence of stochastic gradient descent methods without strong-convexity assumptions.

2 Policy evaluation with the variance-reduction term

We recall that this work focuses on the policy evaluation process. Therefore, u is a fixed control function, and $v(\cdot, \theta)$ is the learnt value function approximating V^u . From now on, for simplicity of the notations, we assume that σ is a constant positive real number. It is straightforward to extend all the results proved here to the case where we take $\tilde{\sigma}$ a matrix-valued function such that $\tilde{\sigma}^\top \tilde{\sigma} \geq \sigma^2 I_d$.

2.1 Short review of the continuous problem and boundary conditions

The value function V^u , defined in (1.3), satisfies the following partial differential equation (Bellman, 1966)

$$(2.1) \quad \mathcal{L}V^u(x) = \rho V^u(x) - \frac{\sigma^2}{2} \Delta_x V^u(x) - \nabla_x V^u(x) \cdot b(x, u(x)) - r(x, u(x)) = 0.$$

Up to now we intentionally omitted to mention the boundary conditions on the state space.

In the following, we will make the simplifying assumption that the state space is the d -dimensional torus, defined as the hypercube $[0, 1]^d$ with periodic boundary conditions, i.e., $\Omega = \mathbb{T}^d = \mathbb{R}^d / \mathbb{Z}^d$. Indeed, the choice of the boundary conditions is in general a difficult problem for continuous state settings, especially when considering PDE-based model, like here with (2.1). More precisely, here, we need them to allow the existence of stationary measures for the continuous and discrete dynamics, and we need a priori estimates on the continuous stationary measure m (more precisely, we need a uniform bound on $\nabla_x \ln m$, see Lemma B.3 in the Appendix). In an attempt to separate difficulties, and to stay focused on the main ideas, we prefer to only consider the torus, where these requirements are easily met, rather than overwhelming the readers with unnecessary technical difficulties due to boundary conditions. See Remark 2.2 below, for some insights on how to extend our results to other domains.

Classical results on second-order elliptic equations imply that if we find a candidate function V such that $\mathcal{L}V$ is small, then V and V^u are close to each other. For instance, see the following lemma.

Lemma 2.1. *Assume that r and b are uniformly bounded on the graph of u . There exists a constant $C > 0$ such that if $V \in H^2(\Omega; \mathbb{R})$, then we have*

$$\|V - V^u\|_{H^2} \leq C \|\mathcal{L}V\|_\infty.$$

The proof consists in: first, applying the maximum principle (see Evans (2010), Section 6.4, for instance) to get $\|V - V^u\|_\infty \leq \rho^{-1} \|\mathcal{L}V\|_\infty$; second, using Theorem 8.12 from Gilbarg and Trudinger

(2001) which yields the desired inequality. More elliptic results, either on weak or strong derivatives, are presented by Ladyzhenskaia and Uraltseva (1968); Krylov (1996).

Regarding Lemma 2.1, an alternative, that will not be investigated here, is to try to minimize the residual of $\mathcal{L}(\cdot, \theta)$ by learning (Lagaris et al., 1998; Rudd, 2013; Sirignano and Spiliopoulos, 2018). Here, we prefer to use the temporal difference instead, in order to make links with usual reinforcement learning algorithms, and to avoid computing second-order derivatives which may be large for high-frequency parametrizations.

Let us mention for later use, that the stationary law $m \in \mathcal{P}(\mathbb{R}^d)$ of the continuous dynamics (1.2) satisfies the following PDE,

$$(2.2) \quad -\frac{\sigma^2}{2}\Delta_x m - \operatorname{div}(b(x, u(x))m(x)) = 0.$$

This may be seen as the dual equation to the homogeneous counterpart of (2.1).

Remark 2.2. *Let us give some hints to extend the present results to other state spaces or boundary conditions. The case $\Omega = \mathbb{R}^d$ is more involved than the torus, but a simple case, to get some insights, consists in taking a drift function deriving from a potential, i.e., $b(x, u(x)) = \nabla_x U(x)$ for some U . In this case, there exists a stationary probability measure of the continuous dynamics if $Z = \int_{\Omega} e^{-2U(x)/\sigma^2} dx$ is finite, given by $m = e^{-2U/\sigma^2}/Z$. Moreover, here, $\nabla_x \ln(m)$ is bounded if and only if $\nabla_x U$ is bounded. This simple example emphasizes the fact that, for \mathbb{R}^d , some restrictions should be satisfied for the results to hold, for instance that the drift function should be bounded and pointing out in the direction of a compact subset of \mathbb{R}^d , with sufficient magnitude.*

Alternatively, one may be interested in considering Ω as a smooth bounded subset of \mathbb{R}^d with, for instance, Dirichlet or Neumann conditions (Evans, 2010). In this case, other restrictions appear, but we also believe that our results may be adapted up to making additional assumptions.

2.2 First results on the asymptotic behavior of standard algorithms

Let us start with the residual gradient algorithm (Baird, 1995). It consists in learning to minimize the mean of the square of the temporal difference by a standard stochastic gradient descent method. The variance of the latter quantity can be decomposed in two terms in the following way,

$$(2.3) \quad \mathbb{E}_{(X, \xi)} \left[\left| \tilde{\delta}_{\Delta t}(X, \xi, \theta) \right|^2 \right] = \underbrace{\mathbb{E}_X \left[\mathbb{E}_{\xi} \left[\tilde{\delta}_{\Delta t}(X, \xi, \theta) \middle| X \right]^2 \right]}_{\text{Bellman error}} + \underbrace{\mathbb{E}_X \left[\operatorname{Var}_{\xi} \left(\tilde{\delta}_{\Delta t}(X, \xi, \theta) \middle| X \right) \right]}_{\text{perturbating term}}.$$

A similar formula holds when $\tilde{\delta}$ is replaced by δ . The asymptotic behavior of the terms inside the expectation \mathbb{E}_X , at fixed x , is characterized in the following lemma.

Lemma 2.3. *Assume that r and b are uniformly bounded on the graph of u , and that v admits bounded continuous derivatives in x everywhere up to order two. Let ξ be normally distributed with mean zero and covariance identity. The means and variances of $\tilde{\delta}$ and δ satisfy*

$$\lim_{\Delta t \rightarrow 0} \mathbb{E}_{\xi} [\tilde{\delta}_{\Delta t}(x, \xi, \theta)] = \mathcal{L}v(x, \theta), \quad \text{and} \quad \lim_{\Delta t \rightarrow 0} \Delta t \operatorname{Var}_{\xi}(\tilde{\delta}_{\Delta t}(x, \xi, \theta)) = \sigma^2 |\nabla_x v(x, \theta)|^2,$$

$$\lim_{\Delta t \rightarrow 0} \mathbb{E}_{\xi} [\delta_{\Delta t}(x, \xi, \theta)] = \mathcal{L}v(x, \theta), \quad \text{and} \quad \lim_{\Delta t \rightarrow 0} \operatorname{Var}_{\xi}(\delta_{\Delta t}(x, \xi, \theta)) = \frac{\sigma^4}{2} \operatorname{tr}(D_x^2 v(x, \theta)^2).$$

The proof is straightforward, using a similar expansion as the one proved in Lemma B.1, below in the Appendix, but only up to order two instead of order four, and Lemma B.6, also in the Appendix.

Passing to the limit $\Delta t \rightarrow 0$ in (2.3), we obtain

$$(2.4) \quad \lim_{\Delta t \rightarrow 0} \mathbb{E}_{(X, \xi)} \left[\left| \tilde{\delta}_{\Delta t}(X, \xi, \theta) \right|^2 \right] = \begin{cases} +\infty & \text{if } v(\cdot, \theta) \text{ is not constant,} \\ \mathbb{E} \left[(\rho C + r(X, u(X)))^2 \right] & \text{if } v(\cdot, \theta) = C. \end{cases}$$

The latter formula yields our first contribution for the residual gradient: when Δt tends to zero, the perturbing term in the decomposition (2.3) totally overwhelms the Bellman error and the only approximations that can be learnt are constant functions.

The same arguments used on δ implies,

$$(2.5) \quad \lim_{\Delta t \rightarrow 0} \mathbb{E}_{(X, \xi)} \left[|\delta_{\Delta t}(X, \xi, \theta)|^2 \right] = \mathbb{E}_X [\mathcal{L}v(X, \theta)^2] + \frac{\sigma^4}{2} \mathbb{E}_X [\text{tr}(D_x^2 v(X, \theta)^2)].$$

In this case, learning is possible even if we see an additional term appearing on the right-hand side. Comments on this additional term and its regularizing effect, and alternatives, are presented in Section 2.3.

Now let us pass to TD(0); it consists in making step in the direction $\nabla_{\theta} v$ with $\tilde{\delta}$ as the probing signal, i.e.,

$$\theta \leftarrow \theta + \gamma \tilde{\delta}_{\Delta t}(X, \xi, \theta) \nabla_{\theta} v(X, \theta),$$

where $\gamma > 0$ is the learning rate. Similar decomposition and analysis of the asymptotic behavior as in (2.3) and Lemma 2.3 respectively, yield

$$\text{Var} \left(\tilde{\delta}_{\Delta t} \nabla_{\theta} v \right) \sim \text{Var} (\mathcal{L}v \nabla_{\theta} v) + \frac{\sigma^2}{2\Delta t} \mathbb{E} \left[|\nabla_x v|^2 |\nabla_{\theta} v|^2 \right],$$

where the arguments of v and $\mathcal{L}v$ are (X, θ) and the ones of $\tilde{\delta}$ are (X, ξ, θ) . This implies the rest of our first contribution, for the case of TD(0). Indeed, learning is impossible at the limit $\Delta t = 0$ because the variance of the update becomes infinite (except if $|\nabla_x v| |\nabla_{\theta} v|$ is uniformly equal to zero which does not seem likely to happen for a large class of functions v).

Repeating the same arguments for δ instead of $\tilde{\delta}$, we obtain

$$\lim_{\Delta t \rightarrow 0} \text{Var} (\delta_{\Delta t} \nabla_{\theta} v) = \text{Var} (\mathcal{L}v \nabla_{\theta} v) + \frac{\sigma^4}{2} \mathbb{E} \left[\text{tr}((D_x^2 v)^2) |\nabla_{\theta} v|^2 \right].$$

Therefore, the latter variance stays uniformly bounded when Δt tends to zero and learning remains possible even at the limit.

2.3 Description of algorithms

To define our learning algorithms, we assume that we access a sequence of observations $(X_k, \xi_k)_{k \geq 0}$. Those observations are assumed to be independent, moreover for $k \geq 0$, ξ_k is normally distributed with mean zero and covariance identity, and X_k is distributed according to m_k , the stationary distribution of the discrete dynamics (1.4) with time step $\Delta t_k > 0$. The sequence (Δt_k) is assumed to be nonincreasing and convergent to zero. In a sense that will be made clear in Theorem 3.1 below, $(m_k)_{k \geq 0}$ is weakly convergent to m , the stationary distribution of the continuous dynamics satisfying (2.2).

In the following, we will have some use of the limiting stochastic processes, therefore we introduce X and ξ two random variables, independent from all the random variables used in the learning algorithms, with laws m satisfying (2.2) and $\mathcal{N}(0, I_d)$ respectively.

The TD(0) algorithm is defined through the following sequence of parameters

$$(TD0) \quad \theta_{k+1} = \theta_k - \gamma_k (\delta_k \varphi(X_k) + \mu \theta_k), \quad \text{with } \delta_k = \delta_{\Delta t_k}(X_k, \xi_k, \theta_k),$$

where $\mu \geq 0$ is a relaxation parameter. Then, we define the residual gradient algorithm by the following sequence,

$$(RG) \quad \theta_{k+1} = \theta_k - \gamma_k (\nabla_{\theta} |\delta_k|^2 + \mu \theta_k), \quad \text{with } \delta_k = \delta_{\Delta t_k}(X_k, \xi_k, \theta_k).$$

When $\mu = 0$, the latter algorithm can be viewed as a stochastic gradient descent method to reduce the right-hand side of (2.5). This right-hand side consists in two terms, the first one being the Bellman error and the second one being an additional term. This additional term penalizes functions with high derivative norms, consequently, it may be viewed as a regularizing term. In some sense, this has a similar effect as taking $\mu > 0$. This motivates our interest in Algorithm (RG) even if it should not converge to the minimizer of the Bellman error even when $\mu = 0$.

Let us propose another algorithm, also based on residual gradient, which will converge to the minimizer of the Bellman error, without additional term. More precisely, this third algorithm is the multiple-step counterpart to (RG), i.e.,

$$(MS-RG) \quad \theta_{k+1} = \theta_k - \gamma_k (\nabla_{\theta} |\delta_k|^2 + \mu \theta_k), \quad \text{with } \delta_k = \frac{1}{n_k} \sum_{i=0}^{n_k-1} \delta_{\Delta t_k}(X_{k,t_i}, \xi_{k,i}, \theta_k),$$

where $X_{k,0} = X_k$ and $X_{k,t_{i+1}} = S_{\Delta t_k}(X_{k,t_i}, \xi_{k,i})$, for $0 \leq i < n_k$ and $n_k \geq 1$ a sequence converging to infinity. Three other alternatives converging to the minimizer of the Bellman error are presented in Section 4.1.

Let us define F as the criterion to minimize when using (RG) or (MS-RG), by,

$$(2.6) \quad F(\theta) = \begin{cases} \mathbb{E} [|\mathcal{L}v(X, \theta)|^2] + \frac{\sigma^4}{2} \mathbb{E} [\text{tr}(D_x^2 v(X, \theta)^2)] + \frac{\mu}{2} |\theta|^2 & \text{for (RG),} \\ \mathbb{E} [|\mathcal{L}v(X, \theta)|^2] + \frac{\mu}{2} |\theta|^2 & \text{for (MS-RG).} \end{cases}$$

Therefore the natural candidate for the convergence of the residual gradient algorithms is $\theta_{RG}^* = \text{argmin}_{\theta \in \Theta} F(\theta)$. In particular, using $F(\theta_{RG}^*) \leq F(0)$, we obtain $\frac{\mu}{2} |\theta_{RG}^*|^2 \leq \mathbb{E} [|r(X)|^2]$.

Concerning Algorithm (TD0), the natural candidate for convergence is θ_{TD}^* satisfying:

$$(2.7) \quad \mathbb{E} [\varphi \mathcal{L}v(X, \theta_{TD}^*)] = -\mu \theta_{TD}^*.$$

This may be rewritten $(\mu I_d + H)\theta_{TD}^* + b = 0$, with $H = \mathbb{E} [\varphi(X) \bar{\mathcal{L}} \varphi(X)^{\top}]$ and $b = \mathbb{E} [r(X) \varphi(X)]$. Using the fact that the symmetric part of H is positive (see Lemma B.3 in the Appendix below), we get $\mu |\theta_{TD}^*|^2 \leq (\theta_{TD}^*)^{\top} (\mu I_d + H) \theta_{TD}^* \leq |b| |\theta_{TD}^*|$, which implies $|\theta_{TD}^*| \leq \mu^{-1} |\mathbb{E} [r(X) \varphi(X)]|$.

Let us introduce the projected counterpart of the latter three algorithms,

$$(P-TD0) \quad \theta_{k+1} = \Pi_B (\theta_k - \gamma_k (\delta_k \varphi(X_k) + \mu \theta_k)), \quad \text{with } \delta_k = \delta_{\Delta t_k}(X_k, \xi_k, \theta_k),$$

$$(P-RG) \quad \theta_{k+1} = \Pi_B (\theta_k - \gamma_k (\nabla_{\theta} |\delta_k|^2 + \mu \theta_k)), \quad \text{with } \delta_k = \delta_{\Delta t_k}(X_k, \xi_k, \theta_k),$$

$$(P-MS-RG) \quad \theta_{k+1} = \Pi_B (\theta_k - \gamma_k (\nabla_{\theta} |\delta_k|^2 + \mu \theta_k)), \quad \text{with } \delta_k = \frac{1}{n_k} \sum_{i=0}^{n_k-1} \delta_{\Delta t_k}(X_{k,t_i}, \xi_{k,i}, \theta_k),$$

where B is a compact ball of \mathbb{R}^d centered at zero and containing θ_{TD}^* or θ_{RG}^* according to the chosen algorithm (for instance, the radius may be $\mu^{-1} |\mathbb{E} [r(X) \varphi(X)]|$ for (P-TD0), or $\sqrt{2\mu^{-1} \mathbb{E} [|r(X)|^2]}$ for (P-RG) and (P-MS-RG)).

3 Main results

3.1 Assumptions

Since the control u is fixed, in this section, we omit the dependence of r and b in their second argument. We say that a function is C^{ℓ} for $\ell \geq 1$ if it admits continuous and bounded derivatives up to order ℓ . Let us make the following assumptions:

H1 The function v is linear with respect to θ , more precisely the set of parametrized functions is $\mathcal{V}_\theta = \{x \mapsto v(x, \theta) = \theta^\top \varphi(x), \theta \in \Theta\}$ for $\varphi : \Omega \rightarrow \mathbb{R}^{d_\theta}$, and $\Theta = \mathbb{R}^{d_\theta}$ for some $d_\theta \geq 1$.

H2 The functions r and b are C^4 .

H3 The feature vector φ is C^6 .

The first assumption, **H1**, is common in the theoretical literature (Polyak and Juditsky, 1992; Tsitsiklis and Van Roy, 1996; Bach and Moulines, 2013), since very little is known about theoretical guarantees in the non-linear case. It is also common to assume some regularity on both the functions of the model, i.e., b and r , and the feature vector φ , like in Assumptions **H2** and **H3**.

Here, the highest needed orders of regularity come from the fact that the observations $(X_k)_{k \geq 0}$ are not distributed according to m the stationary measure of the continuous dynamics (1.2). If they were, **H2** and **H3** could be relaxed: it would be sufficient to assume that r and b are only bounded, and that φ is C^4 (so that we can still use Lemma B.1 from the appendix).

However, here, we assume that $(X_k)_{k \geq 0}$ are distributed through $(m_k)_{k \geq 0}$, the stationary measures of its discrete dynamics (1.4) with $(\Delta t_k)_{k \geq 0}$ as time step, respectively. Consequently, we need to apply the following theorem on the weak convergence of $(m_k)_{k \geq 0}$, with functions r , b , v and its derivatives up to order two. This motivates Assumptions **H2** and **H3**.

Theorem 3.1 (Theorem 14.5.1 from Kloeden and Platen (1992)). *For $f \in C^4(\Omega; \mathbb{R})$, there exists $C > 0$ depending only on the C^4 -norm of f such that $|\mathbb{E}[f(X_k) - f(X)]| \leq C \Delta t_k$.*

3.2 Simple convergence results with relaxed algorithms and decreasing learning steps

This section may be thought of as a warm-up for the next one. We prove convergence of the projected algorithms with $\mu > 0$. This allows to keep the proofs simple and easy to compare with the literature. In particular, we use the common decreasing assumption on the learning step, i.e., that it is proportional to $1/(\mu(k+1))$, and we obtain the usual convergence rate in $1/k$.

Depending on which algorithm is used, we assume that there exists $c > 0$ such that,

$$\mathbf{H4(P-TD0)} \quad \mu > 0, \gamma_k = \frac{2}{\mu(k+1)} \text{ and } \Delta t_k \leq c/\sqrt{k+1}, \text{ for } k \geq 0.$$

$$\mathbf{H4(P-RG)} \quad \mu > 0, \gamma_k = \frac{4}{\mu(k+1)} \text{ and } \Delta t_k \leq c/\sqrt{k+1}, \text{ for } k \geq 0.$$

$$\mathbf{H4(P-MS-RG)} \quad \mu > 0, \gamma_k = \frac{4}{\mu(k+1)}, n_k \geq c^{-1}\sqrt{k+1} \text{ and } n_k \Delta t_k \leq c/\sqrt{k+1}, \text{ for } k \geq 0.$$

Theorem 3.2. *Assume **H1**, **H2**, **H3** and **H4(P-TD0)** or **H4(P-RG)** or **H4(P-MS-RG)**. The sequence $(\theta_k)_{k \geq 0}$ is convergent, and there exists $C > 0$ such that, for $k \geq 1$,*

$$\begin{aligned} |\theta_k - \theta_{TD}^*|^2 &\leq (1 + \mu^{-2}) \frac{C}{\mu^2 k} \\ |\theta_k - \theta_{RG}^*|^2 &\leq (1 + \mu^{-1}) \frac{C}{\mu^2 k} \end{aligned}$$

where the first line corresponds to $TD(0)$ and the second to residual gradient methods. Moreover, if we assume that $|\theta_{TD}^*| \leq C$ (resp. $|\theta_{RG}^*| \leq C$), then we can remove the term $(1 + \mu^{-2})$ (resp. $(1 + \mu^{-1})$) in the first (resp. second) line.

The factor $1/\mu^2$ in the latter theorem is a consequence of the decreasing learning steps being proportional to $1/\mu$. Under an averaging method, with the same sequence of learning steps, we would still get a factor $1/\mu$. This motivates the next section, where we still get convergence even with $\mu = 0$.

3.3 Averaging method with constant learning step and no relaxation

In this section, we use the Polyak-Juditsky averaging method, see Polyak and Juditsky (1992), to accelerate the convergence of the TD(0) algorithm. In the same spirit as the results from Bach and Moulines (2013), we obtain optimal convergence speed with constant learning step and without relaxation assumption, i.e., $\mu = 0$ and no projection map used. Therefore, following (2.7), the potential limit θ_{TD}^* satisfies

$$(3.1) \quad \Pi \mathcal{L}v(\cdot, \theta_{TD}^*) = 0,$$

where Π is the $L^2(m)$ -orthogonal projector onto \mathcal{V}_Θ .

Theorem 3.3. *Assume **H1**, **H2** and **H3**. and that θ_{TD}^* is bounded. If $\sum_{i=0}^{\infty} \Delta t_i^2$ is finite, there exist $C, R > 0$ such that, the following inequality holds for $\gamma < R^{-2}$, $k \geq 1$,*

$$\begin{aligned} \mathbb{E} [v(X, \bar{\theta}_k) - v(X, \theta_{TD}^*)] &\leq \frac{C}{\gamma k} + \frac{C(d + \text{tr}(HH^{-\top}))}{k}, \\ \mathbb{E} [v(X, \bar{\theta}_k) - v(X, \theta_{RG}^*)] &\leq \frac{C}{\gamma k} + \frac{Cd}{k}, \end{aligned}$$

where the first inequality happens with TD(0) and the second with the residual gradient methods, $\bar{\theta}_k = \frac{1}{k} \sum_{i=0}^{k-1} \theta_i$, for $k \geq 1$, and $H = \mathbb{E} [\varphi(X) \bar{\mathcal{L}}\varphi(X)^\top]$.

If instead we assume that $\sum_{i=0}^{k-1} \Delta t_i^2 \leq a \ln(1+k)$ for some $a > 0$ for any $k \geq 0$, then for any $\varepsilon > 0$ there exists $C, R > 0$ such that for $\gamma < R^{-2}$, $k \geq 0$, the latter inequalities are replaced with the following ones respectively

$$\begin{aligned} \mathbb{E} [v(X, \bar{\theta}_k) - v(X, \theta_{TD}^*)] &\leq \frac{C}{\gamma k} + \frac{C(d + \text{tr}(HH^{-\top}))}{k^{1-\varepsilon}}, \\ \mathbb{E} [v(X, \bar{\theta}_k) - v(X, \theta_{TD}^*)] &\leq \frac{C}{\gamma k} + \frac{Cd}{k^{1-\varepsilon}}. \end{aligned}$$

The proof is adapted from Bach and Moulines (2013) with the extra difficulties that the linear operator applied to θ_k in (TD0), (RG) and (MS-RG) are different for any $k \geq 0$, it is not symmetric in the case of (TD0), and its symmetric part has no interesting properties (only the symmetric part of the expectation of its limit when $k \rightarrow \infty$ have useful properties). Moreover, our sequence of stochastic estimators have vanishing biases that introduce new terms to bound in the proof, this leads to the necessity to add assumptions on $\sum_i \Delta t_i^2$.

4 Possible extensions

4.1 Alternatives to the multiple-step residual gradient method

In this section, we present three residual gradient algorithms that are alternatives to the multiple-step method (MS-RG) in the sense that they converge to the same minimizer as (MS-RG). In particular, Theorems 3.2 and 3.3 hold with this three alternatives.

Vanishing viscosities. This algorithm consists in the following induction relation,

$$(\sigma\text{RG}) \quad \theta_{k+1} = \theta_k - \gamma_k (\nabla_\theta |\delta_k|^2 + \mu \theta_k), \quad \text{with } \delta_k = \delta_{\sigma_k, \Delta t_k}(X_k, \xi_k, \theta_k),$$

where $\delta_{\sigma_k, \Delta t_k}$ is $\delta_{\Delta t_k}$ where σ has been replaced by σ_k in (1.7) and (1.4). Here, we assume that we may choose the intensity of the noise, i.e., σ_k , which is only possible for Class **C2**. The conclusions of Theorem 3.2 then hold with $\gamma_k = \frac{4}{\mu(k+1)}$, $\Delta t_k \leq c/\sqrt{k+1}$, and $\sigma_k \leq ck^{-\frac{1}{8}}$, for $k \geq 0$.

Using mini-batches. Another alternative consists in using mini-batches, i.e.,

$$(MB\text{-RG}) \quad \theta_{k+1} = \theta_k - \gamma_k (\nabla_{\theta} |\delta_k|^2 + \mu \theta_k), \quad \text{with } \delta_k = \frac{1}{N_k} \sum_{i=1}^{N_k} \delta_{\Delta t_k}(X_k, \xi_{k,i}, \theta_k),$$

where $(N_k)_{k \geq 0}$ are the size of the mini-batches. The conclusions of Theorem 3.2 then hold with $\gamma_k = \frac{4}{\mu(k+1)}$, $\Delta t_k \leq c/\sqrt{k+1}$, and $N_k \geq c^{-1}\sqrt{k}$, for $k \geq 0$.

Changing the law of the noise. Note that the perturbing term from (2.5) comes from the variance of a term involving $\xi_k^{\top} D_x^2 v \xi_k - \Delta_x v$. Let us make the simple observation that, in dimension $d = 1$, the latter expression is null if ξ_k is a Rademacher random variable. This argument can be generalized to dimension $d > 1$. Since $D^2 v(X_k, \theta_k)$ is symmetric, we can find P an orthogonal matrix and D a diagonal matrix such that $D_x^2 v(X_k, \theta_k) = P^{\top} D P$. Therefore, it we can take $\xi_k = P^{\top} \zeta_k$ where ζ_k is a random vector, each of its coordinate being an independent Rademacher random variable.

Using Donsker's theorem (Donsker, 1951), the random process at the limit is still a Brownian motion even if the increments before convergence are not Gaussian anymore. However, the weak convergence of the sequence $(m_k)_{k \geq 0}$ is slower here: Δt_k is replaced by $\Delta t_k^{\frac{1}{2}}$ (this is a consequence of the central limit theorem). The conclusions of Theorem 3.2 then hold with $\gamma_k = \frac{4}{\mu(k+1)}$ and $\Delta t_k \leq c/(k+1)$, for $k \geq 0$.

4.2 Non-vanishing time steps

In practice, we never iterate until convergence, and the algorithms are stopped after a certain number of iterations are computed, or after some threshold is reached. Therefore, we never really consider the case $\Delta t \rightarrow 0$.

Consequently, an interesting different setting consists in considering non-vanishing time steps. In this section, we consider $(\Delta t_k)_{k \geq 0}$ to be constant and equal to some $\Delta t > 0$. Classical results of convergence of TD(0) and residual gradient methods apply here to state similar results as the one proved in the present work. However, the constants in the rates of convergence of the latter classical results will depend on Δt . Typically, those constants will tend to infinity when Δt tends to zero.

Here, we can do better by extending the results of the vanishing setting to the non-vanishing setting. Namely, if Δt is small enough, i.e., $\Delta t \leq \Delta t_0$ for some $\Delta t_0 > 0$, the results of Sections 3.2 and 3.3 hold with constants independent of Δt . This extension is justified because all the inequalities and arguments needed to prove the original results, hold for small Δt up to changing a little bit the associated constants, using Theorem 3.1.

5 Discussion

In the present work, we proved that standard reinforcement learning method based on the temporal difference are not adapted to solve continuous stochastic optimization, nor their discretizations using small time-steps. We proposed a correction to the temporal difference, in order to overcome the latter problem and obtain robust algorithms with respect to vanishing time steps. We proved two types of convergence results of our algorithms. The first ones consisting in the counterparts of simple standard convergence results, they are easy to compare with the literature. The second ones, might be seen as the analogous to some state-of-the-art convergence results of stochastic gradient descent without strong-convexity assumption.

After this work on the policy evaluation process, the natural next step is to make an analysis of the policy improvement process, in order to get entire algorithms converging to an approximation of the minimizer of a stochastic continuous optimization problem.

Furthermore, we would like to study the model-free counterparts of the algorithms studied here. Indeed, our algorithms rely on the assumption that we can observe the noise of the dynamics (or at least the drift), they are model-based, which may be limiting for some models in reinforcement learning. Making our algorithms be model-free is possible by learning the correction in the same time as the value function.

Other extensions would be worthy of future works. The convergence rates in Section 3.2 are independent of the dimension of the parametrized space, therefore, they might be extended to infinite dimensions. These results still rely on the linearity assumption, so natural candidates that we would like to investigate are reproducing kernel Hilbert spaces.

Acknowledgements. We thank Justin Carpentier for fruitful discussions related to this work. This work was funded in part by the French government under management of Agence Nationale de la Recherche as part of the “Investissements d’avenir” program, reference ANR-19-P3IA-0001(PRAIRIE 3IA Institute). We also acknowledge support from the European Research Council (grant SEQUOIA 724063).

References

- Francis Bach and Eric Moulines. Non-strongly-convex smooth stochastic approximation with convergence rate $O(1/n)$. *Advances in Neural Information Processing Systems*, 26:773–781, 2013.
- Leemon Baird. Residual algorithms: Reinforcement learning with function approximation. In *Proceedings of the International Conference on Machine Learning*, pages 30–37, 1995.
- Leemon C. Baird III. Advantage updating. Technical report, Wright Lab Wright-Patterson AFB OH, 1993.
- Richard Bellman. Dynamic programming. *Science*, 153(3731):34–37, 1966.
- Steven Bradtke and Michael Duff. Reinforcement learning methods for continuous-time Markov decision problems. *Advances in Neural Information Processing Systems*, 7, 1994.
- Peter Dayan and Satinder P. Singh. Improving policies without measuring merits. *Advances in neural information processing systems*, pages 1059–1065, 1996.
- Monroe D. Donsker. An invariance principle for certain probability limit theorems. AMS, 1951.
- Kenji Doya. Reinforcement learning in continuous time and space. *Neural Computation*, 12(1): 219–245, 2000.
- Lawrence C Evans. *Partial differential equations*, volume 19. American Mathematical Soc., 2010.
- David Gilbarg and Neil S. Trudinger. *Elliptic partial differential equations of second order*. Classics in Mathematics. Springer-Verlag, Berlin, 2001. ISBN 3-540-41160-7. Reprint of the 1998 edition.
- Donald E. Kirk. Optimal control theory: An introduction. 1970.
- Peter E. Kloeden and Eckhard Platen. *Numerical Solution of Stochastic Differential Equations*, volume 23 of *Applications of Mathematics (New York)*. Springer-Verlag, Berlin, 1992.
- Vijay R. Konda and John N. Tsitsiklis. Actor-critic algorithms. In *Advances in Neural Information Processing Systems*, pages 1008–1014, 2000.

- Nikolaï Vladimirovich Krylov. *Lectures on elliptic and parabolic equations in Holder spaces*. Number 12. American Mathematical Soc., 1996.
- Olga A. Ladyzhenskaia and Nina N. Uraltseva. *Linear and quasilinear elliptic equations [by] Olga A. Ladyzhenskaya and Nina N. Uraltseva. Translated by Scripta Technica. Translation editor: Leon Ehrenpreis*. Academic Press New York, 1968.
- Isaac E. Lagaris, Aristidis Likas, and Dimitrios I. Fotiadis. Artificial neural networks for solving ordinary and partial differential equations. *IEEE Transactions on Neural Networks*, 9(5): 987–1000, 1998.
- Quentin Le Lidec, Louis Montaut, Cordelia Schmid, Ivan Laptev, and Justin Carpentier. Leveraging randomized smoothing for optimal control of nonsmooth dynamical systems. 2021.
- Michael Lutter, Shie Mannor, Jan Peters, Dieter Fox, and Animesh Garg. Value iteration in continuous actions, states and time. *arXiv preprint arXiv:2105.04682*, 2021.
- Huyen Pham, Xavier Warin, and Maximilien Germain. Neural networks-based backward scheme for fully nonlinear PDEs. *SN Partial Differential Equations and Applications*, 2(1):1–24, 2021.
- Boris Polyak and Anatoli B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30:838–855, 1992.
- Keith Rudd. *Solving partial differential equations using artificial neural networks*. PhD thesis, Duke University, 2013.
- Justin Sirignano and Konstantinos Spiliopoulos. DGM: A deep learning algorithm for solving partial differential equations. *Journal of computational physics*, 375:1339–1364, 2018.
- Richard S. Sutton. Learning to predict by the methods of temporal differences. *Machine Learning*, 3(1):9–44, 1988.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: an Introduction*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA, second edition, 2018.
- Corentin Tallec, Léonard Blier, and Yann Ollivier. Making deep Q-learning methods robust to time discretization. In *International Conference on Machine Learning*, pages 6096–6104. PMLR, 2019.
- John Tsitsiklis and Benjamin Van Roy. Analysis of temporal-difference learning with function approximation. *Advances in Neural information Processing Systems*, 9, 1996.
- Haoran Wang, Thaleia Zariphopoulou, and Xun Yu Zhou. Reinforcement learning in continuous time and space: A stochastic control approach. *Journal of Machine Learning Research*, 21: 198–1, 2020.
- Cagatay Yildiz, Markus Heinonen, and Harri Lähdesmäki. Continuous-time model-based reinforcement learning. In *International Conference on Machine Learning*, pages 12009–12018. PMLR, 2021.

A Proof of the main results

Here, C is a constant that can change from line to line and is independent from $(\gamma_k)_{k \geq 0}$, $(\theta_k)_{k \geq 0}$ and μ .

A.1 Proof of Theorem 3.2 with residual gradient methods

Let us start by proving the following theorem on stochastic gradient descent methods.

Theorem A.1. *Let $f : \Theta \rightarrow \mathbb{R}$ be μ -convex, L -semi-concave, and such that $\theta^* = \operatorname{argmin}_\theta f(\theta)$ satisfies $|\theta^*| \leq M$ for some $M > 0$. For $\theta_0 \in \Theta$, the sequence $(\theta_k)_{k \geq 0}$ is defined by induction using the following projected stochastic gradient descent method,*

$$\theta_{k+1} = \Pi_{B(0, M)}(\theta_k - \gamma_k g_k),$$

for $k \geq 0$, where $\gamma_k > 0$ is convergent to zero, and $\sum_{k \geq 0} \gamma_k = \infty$, $|\mathbb{E}[g_k | \theta_k] - f'(\theta_k)| \leq (1 + |\theta_k|)\varepsilon_k$, $\varepsilon_k \in \mathbb{R}_+$ is convergent to zero, and $\mathbb{E}[|g_k|^2 | \theta_k] \leq C(1 + |\theta_k|^2)$. Then $(\theta_k)_{k \geq 0}$ is convergent in expectation to θ^* , and

$$\mathbb{E}[|\theta_k - \theta^*|^2] \leq 4M^2 e^{-\frac{\mu}{2} \sum_{i=0}^{k-1} \gamma_i} + C(1 + M^2) \sum_{i=0}^{k-1} \gamma_i (\gamma_i + \mu^{-1} \varepsilon_i^2) e^{-\frac{\mu}{2} \sum_{j=i+1}^{k-1} \gamma_j}.$$

Proof. Up to starting the iterative algorithm from θ_1 instead of θ_0 , we may assume that $|\theta_k| \leq M$ for every $k \geq 0$. For $k \geq 0$, let us denote $b_k = |\theta_k - \theta^*|^2$. We recall that $|\Pi_{B(0, M)}(\theta) - \theta^*| \leq |\theta - \theta^*|$ for any $\theta \in \Theta$, since $\theta^* \in B(0, M)$. This and the induction relation satisfied by θ_k , imply

$$\begin{aligned} b_{k+1} &\leq \mathbb{E}[|\theta_k - \theta^* - \gamma_k g_k|^2] \\ &\leq b_k - 2\gamma_k \mathbb{E}[(\theta_k - \theta^*)^\top g_k] + \gamma_k^2 \mathbb{E}[|g_k|^2] \\ &\leq b_k - 2\gamma_k \mathbb{E}[(\theta_k - \theta^*)^\top \mathbb{E}[g_k | \theta_k]] + \gamma_k^2 \mathbb{E}[\mathbb{E}[|g_k|^2 | \theta_k]] \\ &\leq b_k - 2\gamma_k \mathbb{E}[(\theta_k - \theta^*)^\top f'(\theta_k)] + 2\gamma_k \varepsilon_k \mathbb{E}[|\theta_k - \theta^*|(1 + |\theta_k|)] + C\gamma_k^2 \mathbb{E}[(1 + |\theta_k|^2)] \\ &\leq b_k - 2\gamma_k \mathbb{E}[f(\theta^*) - f(\theta_k) - \frac{\mu}{2} |\theta_k - \theta^*|^2] + 2(1 + M)\gamma_k \varepsilon_k \mathbb{E}[|\theta_k - \theta^*|] + C(1 + M^2)\gamma_k^2 \\ &\leq (1 - \mu\gamma_k)b_k + \frac{\mu}{2}\gamma_k \mathbb{E}[|\theta_k - \theta^*|^2] + 4(1 + M^2)\mu^{-1}\gamma_k \varepsilon_k^2 + C(1 + M^2)\gamma_k^2 \\ &\leq (1 - \frac{\mu}{2}\gamma_k)b_k + C(1 + M^2)\gamma_k(\mu^{-1}\varepsilon_k^2 + \gamma_k) \\ &\leq e^{-\frac{\mu}{2}\gamma_k} b_k + C(1 + M^2)\gamma_k(\mu^{-1}\varepsilon_k^2 + \gamma_k), \end{aligned}$$

where we used the μ -strong convexity of f to get to the fifth line, and a Young inequality to obtain the sixth line. Therefore, we obtain

$$b_k \leq e^{-\frac{\mu}{2} \sum_{i=0}^{k-1} \gamma_i} b_0 + C(1 + M^2) \sum_{i=0}^{k-1} \gamma_i (\gamma_i + \mu^{-1} \varepsilon_i^2) e^{-\frac{\mu}{2} \sum_{j=i+1}^{k-1} \gamma_j},$$

which leads to the desired inequality using $b_0 \leq (|\theta_0| + |\theta^*|)^2 \leq 4M^2$. \square

Proof of Theorem 3.2 with the residual gradient methods. We only prove the results for (RG). The proof may be repeated to hold for (MS-RG), using Lemma B.2. It consists in checking that we can apply Theorem A.1, using the following notations,

$$f(\theta) = \mathbb{E}[|\mathcal{L}v(X, \theta)|^2] + \frac{\sigma^4}{2} \mathbb{E}[\operatorname{tr}(D_x^2 v(X, \theta)^2)] + \frac{\mu}{2} |\theta|^2, \quad \text{and } g_k = \nabla_\theta |\delta_k|^2 + \mu \theta_k.$$

Thus, we get,

$$\begin{aligned}\mathbb{E}[g_k|\theta_k] &= \mathbb{E}\left[\nabla_\theta\left|\mathcal{L}(X_k, \theta_k) + R_{0,k}^\top\theta_k + \Delta t_k^{\frac{1}{2}}R_{1,k}^\top\theta_k + \Delta t_k R_{2,k}^\top\theta_k\right|^2\right] + \mu\theta_k \\ &= \mathbb{E}\left[\nabla_\theta\left|\mathcal{L}(X_k, \theta_k)\right|^2\right] + \mathbb{E}\left[\nabla_\theta\left|R_{0,k}^\top\theta_k\right|^2\right] + \mu\theta_k + \Delta t_k\mathbb{E}\left[\nabla_\theta\left|R_{1,k}^\top\theta_k\right|^2\right] + 2\Delta t_k\mathbb{E}\left[\nabla_\theta\left(\delta_k R_{2,k}^\top\theta_k\right)\right].\end{aligned}$$

Then, from Theorem 3.1, we obtain

$$\left|\mathbb{E}\left[\nabla_\theta\left|\mathcal{L}(X_k, \theta_k) + R_{0,k}^\top\theta_k + \Delta t_k^{\frac{1}{2}}R_{1,k}^\top\theta_k + \Delta t_k R_{2,k}^\top\theta_k\right|^2\right] + \mu\theta_k - f'(\theta_k)\right| \leq C(1 + |\theta_k|).$$

This implies that $|\mathbb{E}[g_k|\theta_k] - f'(\theta_k)| \leq C\Delta t_k(1 + |\theta_k|)$. The fact that $\mathbb{E}[|g_k|^2|\theta_k] \leq C(1 + |\theta_k|^2)$ is straightforward. Theorem A.1 and the inequalities $|\theta^*|^2 \leq C\mu^{-1}$ and $\exp\left(-\sum_{j=i+1}^k \frac{1}{j}\right) \leq i/k$ for $k > i \geq 0$, conclude the proof. \square

A.2 Proof of Theorem 3.2 with TD(0)

Theorem A.2. Let $A \in \mathbb{R}^{d \times d}$ be a square matrix such that $A + A^\top \geq 2\mu I_d$ for some $\mu > 0$, and $b, \theta^* \in \mathbb{R}^d$ such that $A\theta^* = b$ and $|\theta^*| \leq M$ for some $M \geq 0$. For $\theta_0 \in \Theta$, the sequence $(\theta_k)_{k \geq 0}$ is defined by induction by,

$$\theta_{k+1} = \Pi_{B(0,M)}(\theta_k - \gamma_k g_k),$$

for $k \geq 0$, where $\gamma_k > 0$ is convergent to zero and $\sum_{k \geq 0} \gamma_k = \infty$, $|\mathbb{E}[g_k|\theta_k] - A\theta_k - b| \leq (1 + |\theta_k|)\varepsilon_k$, $\varepsilon_k \in \mathbb{R}_+$ is convergent to zero, and $\mathbb{E}[|g_k|^2|\theta_k] \leq C(1 + |\theta_k|^2)$. Then $(\theta_k)_{k \geq 0}$ is convergent in expectation to θ^* and

$$\mathbb{E}[|\theta_k - \theta^*|^2] \leq 4M^2 e^{-\mu \sum_{i=0}^{k-1} \gamma_i} + C(1 + M^2) \sum_{i=0}^{k-1} \gamma_i (\gamma_i + \mu^{-1} \varepsilon_i^2) e^{-\mu \sum_{j=i+1}^{k-1} \gamma_j}.$$

Proof. Here, C is a constant that can change from line to line and is independent from $(\gamma_k)_{k \geq 0}$, $(\theta_k)_{k \geq 0}$ and μ . Using similar arguments as in the beginning of the proof of Theorem A.1, with the notation $b_k = |\theta_k - \theta^*|^2$ for $k \geq 0$, we obtain,

$$\begin{aligned}b_{k+1} &\leq b_k - 2\gamma_k \mathbb{E}\left[(\theta_k - \theta^*)^\top \mathbb{E}[g_k|\theta_k]\right] + \gamma_k^2 \mathbb{E}\left[\mathbb{E}[|g_k|^2|\theta_k]\right] \\ &\leq b_k - 2\gamma_k \mathbb{E}\left[(\theta_k - \theta^*)^\top (A\theta_k + b)\right] + 2\gamma_k \varepsilon_k \mathbb{E}[|\theta_k - \theta^*|(1 + |\theta_k|)] + C\gamma_k^2 \mathbb{E}[(1 + |\theta_k|^2)] \\ &\leq b_k - \gamma_k \mathbb{E}\left[(\theta_k - \theta^*)^\top (A + A^\top)(\theta_k - \theta^*)\right] + \mu\gamma_k \mathbb{E}[|\theta_k - \theta^*|^2] + 2(1 + M^2)\mu^{-1}\gamma_k \varepsilon_k^2 + C(1 + M^2)\gamma_k^2 \\ &\leq (1 - \mu\gamma_k)b_k + C(1 + M^2)\gamma_k(\mu^{-1}\varepsilon_k^2 + \gamma_k) \\ &\leq e^{-\mu\gamma_k}b_k + C(1 + M^2)\gamma_k(\mu^{-1}\varepsilon_k^2 + \gamma_k),\end{aligned}$$

where we used a Young inequality to get to the third line. We conclude with similar arguments as in the proof of Theorem A.1. \square

Proof of Theorem 3.2 with the TD(0) method. The proof only consists in checking that we can apply Theorem A.2. Using the same notation as in Theorem A.2, we define,

$$A = \mathbb{E}[\varphi(X)\bar{\mathcal{L}}(X)] + \mu I_d, \quad b = \mathbb{E}[r(X)\varphi(X)], \quad \text{and } g_k = \delta_k \varphi(X_k) + \mu\theta_k.$$

Then, we get

$$\begin{aligned}\mathbb{E}[g_k|\theta_k] &= \mathbb{E}\left[\varphi(X_k)\left(\bar{\mathcal{L}}\varphi(X_k) + R_{0,k} + \Delta t_k^{\frac{1}{2}}R_{1,k} + \Delta t_k R_{2,k}\right)\right]\theta_k + \mu\theta_k + \mathbb{E}[\varphi(X_k)r(X_k)] \\ &= \mathbb{E}[\varphi(X_k)\bar{\mathcal{L}}\varphi(X_k)]\theta_k + \mu\theta_k + \mathbb{E}[\varphi(X_k)r(X_k)] + \Delta t_k\mathbb{E}\left[\varphi(X_k)R_{2,k}^\top\right]\theta_k,\end{aligned}$$

where $R_{0,k} = R_0(X_k, \xi_k)$, $R_{1,k} = R_1(X_k, \xi_k)$ and $R_{2,k} = R_2(\Delta t_k, X_k, \xi_k)$ are given in Lemma B.1. From Theorem 3.1, we get

$$|\mathbb{E}[\varphi(X_k)\bar{\mathcal{L}}\varphi(X_k)] - A| \leq C, \quad \text{and} \quad |\mathbb{E}[\varphi(X_k)r(X_k)] - b| \leq C.$$

Therefore, we obtain $|\mathbb{E}[g_k|\theta_k] - A\theta_k - b| \leq C(1 + |\theta_k|)\Delta t_k$. The fact that $\mathbb{E}[|g_k|^2|\theta_k] \leq C(1 + |\theta_k|^2)$ is straightforward. Finally, $A + A^\top \geq 2\mu I_d$ comes from Lemma B.3. Theorem A.2 and the inequalities $|\theta^*| \leq C\mu^{-1}$ and $\exp\left(-\sum_{j=i+1}^k \frac{1}{j}\right) \leq i/k$ for $k > i \geq 0$, conclude the proof. \square

A.3 Proof of Theorem 3.3

Here we only make the proof for TD(0), the proof for the residual gradient methods being similar and simpler since the counterpart for residual gradient of the matrix H defined below is symmetric, simplifying a lot the equations where H , H^\top and there symmetric part S were involved.

We start with the following definitions,

$$\begin{aligned}S &= \rho\mathbb{E}\left[\varphi(X)\varphi(X)^\top\right] + \frac{\sigma^2}{2}\mathbb{E}\left[D_x\varphi(X)D_x\varphi(X)^\top\right] \\ A &= \mathbb{E}\left[\varphi(X)\left(\frac{\sigma^2}{2}\nabla_x \ln(m) + b\right)D_x\varphi(X)^\top\right] \\ H(x) &= \varphi(x)\bar{\mathcal{L}}\varphi(x)^\top \\ H_k(x) &= H(x) + \mathbb{E}[H(X) - H(X_k)] \\ H &= \mathbb{E}[H(X)].\end{aligned}$$

Proof of Theorem 3.3 for TD(0). Here, $C > 0$ stands for a generic constant which value may change from line to line, it depends on the constants in the assumptions and is independent of k , of the smallest eigenvalue of S and of γ .

Using Lemma B.1, we get

$$\theta_{k+1} = \theta_k - \gamma\varphi(X_k)\left(\bar{\mathcal{L}}\varphi(X_k) + R_0(X_k, \xi_k) + \Delta t_k^{\frac{1}{2}}R_1(X_k, \xi_k) + \Delta t_k R_2(\Delta t_k, X_k, \xi_k)\right)^\top \theta_k - \gamma\varphi(X_k)r(X_k),$$

where $R_0(x, \xi)^\top \theta = \frac{\sigma^2}{2}(\xi^\top D_x^2 v(x, \theta)\xi - \Delta_x v(x, \theta))$, and R_1 and R_2 can be read in Lemma B.1, and we get $\mathbb{E}_\xi[R_0(x, \xi)] = \mathbb{E}[R_1(x, \xi)] = 0$. Take $\eta_k = \theta_k - \theta^*$, it satisfies the following induction relation,

$$\eta_{k+1} = (I_d - \gamma H_k(X_k))\eta_k - \gamma(H_k(X_k)\theta^* + \varphi(X_k)r(X_k)) - \gamma\left(H - \mathbb{E}[H(X_k)] + \Delta t_k\varphi(X_k)R_{2,k}^\top\right)(\eta_k + \theta^*),$$

where $H_k(x) = \varphi(x)(\bar{\mathcal{L}}\varphi(x) + R_{0,k} + \Delta t_k^{\frac{1}{2}}R_{1,k})^\top + H - \mathbb{E}[H(X_k)]$, in particular $\mathbb{E}[H_k(X_k)] = H$. One may easily check that η_k can be rewritten as $\eta_k = \sum_{r=0}^{k-1} \eta_k^r$, where η_k^r is defined by,

$$(A.1) \quad \begin{aligned}\eta_{k+1}^r &= (I_d - \gamma H)\eta_k^r + \chi_k^r + \Delta t_k \psi_k^r, \\ \eta_0^0 &= \eta_0, \quad \eta_0^r = 0 \text{ if } r \geq 1,\end{aligned}$$

where χ_k^r and ψ_k^r are defined by

$$(A.2) \quad \begin{aligned} \chi_k^0 &= \gamma(H - H_k(X_k))\theta^* + \gamma(\mathbb{E}[\varphi(X_k)r(X_k)] - \varphi(X_k)r(X_k)), \\ \psi_k^0 &= \gamma\Delta t_k^{-1}(\mathbb{E}[\varphi(X_k)\mathcal{L}v(X_k, \theta^*)] - \mathbb{E}[\varphi(X)\mathcal{L}v(X, \theta^*)]) - \gamma\varphi(X_k)R_{2,k}^\top\theta^*, \\ \chi_k^{r+1} &= \gamma(H - H_k(X_k))\eta_k^r, \\ \psi_k^{r+1} &= \gamma\left(\Delta t_k^{-1}(\mathbb{E}[H(X_k)] - H) - \varphi(X_k)R_{2,k}^\top\right)\eta_k^r, \end{aligned}$$

where we used that $\mathbb{E}[\varphi(X)\mathcal{L}v(X, \theta^*)] = 0$ to get the second line. One may notice that $\eta_k^r = 0$ if $r \geq k$.

First step: getting bounds on the covariance matrices of χ_r^k and ψ_r^k . Here, we prove by induction on r and k that

$$\begin{aligned} \mathbb{E}[\eta_k^r \otimes \eta_k^r] &\leq 3C_k\gamma^r R^{2r} I_d, \\ \mathbb{E}[\chi_k^r \otimes \chi_k^r] &\leq C_k\gamma^{\max(r+1,2)} R^{2r} S, \\ \mathbb{E}[\psi_k^r \otimes \psi_k^r] &\leq \varepsilon C_k\gamma^{\max(r+1,2)} R^{2r} S, \end{aligned}$$

where $R^2 = 3\tilde{C} \left(\|\bar{\mathcal{L}}\varphi + \mathbb{E}[R_0(\cdot, \xi_0)]\|_\infty + \Delta t_0^{\frac{1}{2}}\|R_1(\cdot, \xi)\|_\infty + 2\varepsilon^{-1} \sup_{k \geq 0} \|R_2(\Delta t_k, \cdot, \xi)\|_\infty + 2\varepsilon^{-1} \right)$, $0 < \varepsilon < \Delta_0^{-2}$ is a constant that will be defined later, \tilde{C} is the constant from Lemma B.5 and $C_k = (|\theta^*|^2 + \eta_0^\top S \eta_0) \exp(\varepsilon \sum_{i=0}^{k-1} \Delta t_i^2)$.

For $k \geq 0$, and $r \geq 1$, let us prove the results for $(k+1, r)$ while assuming that it holds for (k, r) , $(k, r-1)$ and $(k+1, r-1)$. For $b_k = \varepsilon\Delta t_k^2$, we get from (A.1) and (B.6),

$$\begin{aligned} \mathbb{E}[\eta_{k+1}^r \otimes \eta_{k+1}^r] &\leq (1 + b_k) \mathbb{E} \left[(I_d - \gamma H) \eta_k^r \otimes \eta_k^r (I_d - \gamma H^\top) \right] + \mathbb{E}[\chi_k^r \otimes \chi_k^r] + \Delta t_k^2 (1 + b_k^{-1}) \mathbb{E}[\psi_k^r \otimes \psi_k^r] \\ &\leq 3C_k\gamma^r R^{2r} (1 + b_k) (I_d - \gamma H)(I_d - \gamma H^\top) + C_k\gamma^{r+1} R^{2r} + \varepsilon C_k \Delta t_k^2 \gamma^{r+1} R^{2r} (1 + b_k^{-1}) \\ &\leq 3C_k\gamma^r R^{2r} (1 + \varepsilon\Delta t_k^2) (I_d - \gamma S) + \gamma^{r+1} R^{2r} C_k (2 + \varepsilon\Delta t_k^2) S \\ &\leq 3C_k\gamma^r R^{2r} (1 + \varepsilon\Delta t_k^2) I_d \leq 3C_k e^{\varepsilon\Delta t_k^2} \gamma^r R^{2r} I_d = 3C_{k+1} \gamma^r R^{2r} I_d. \end{aligned}$$

Then, concerning χ_{k+1}^r , using Lemma B.5, we get

$$\begin{aligned} \mathbb{E}[\chi_{k+1}^r \otimes \chi_{k+1}^r] &\leq 3C_{k+1}\gamma^{r+1} R^{2r-2} \mathbb{E} \left[(H - H_k(X_k))(H - H_k(X_k))^\top \right] \\ &\leq 3C_{k+1}\gamma^{r+1} R^{2r-2} \mathbb{E} \left[H_k(X_k) H_k(X_k)^\top \right] \\ &\leq 3C_{k+1}\gamma^{r+1} R^{2r-2} \left\| \bar{\mathcal{L}}\varphi + \mathbb{E}[R_0(\cdot, \xi_k) + \Delta t_k^{\frac{1}{2}} R_1(\cdot, \xi^k)] \right\|_\infty \mathbb{E} \left[\varphi(X_k) \otimes \varphi(X_k)^\top \right] \\ &\leq C_{k+1}\gamma^{r+1} R^{2r} S. \end{aligned}$$

Finally, using Lemma B.5 once again for ψ_{k+1}^r , we get,

$$\begin{aligned} \mathbb{E}[\psi_{k+1}^r \otimes \psi_{k+1}^r] &\leq 6C_{k+1}\gamma^{r+1} R^{2r-2} \left(\Delta t_k^{-2} (\mathbb{E}[H(X_k)] - H)(\mathbb{E}[H(X_k)] - H)^\top + \mathbb{E}[|R_{2,k}|^2 \varphi(X_k) \otimes \varphi(X_k)] \right) \\ &\leq \varepsilon C_{k+1} \gamma^{r+1} R^{2r} S. \end{aligned}$$

It remains to prove the inequalities for $k = 0$ and $r = 0$. Concerning $r = 0$, the proof is similar but we use the boundedness of θ^* and r instead of the induction assumption. Then $k = 0$ and $r \geq 1$ is straightforward since $\eta_0^r = \chi_0^r = \psi_0^r = 0$.

Second step: getting a bound on $\mathbb{E}[(\bar{\eta}_k^r)^\top S \bar{\eta}_k^r]$. Namely, we will prove that

$$\mathbb{E} \left[(\bar{\eta}_k^r)^\top S \bar{\eta}_k^r \right] \leq \frac{C\gamma^{\max(r-1,0)} R^{2r}}{k} \text{tr}(I_d + H^{-\top} H) \left(\frac{1}{k} \sum_{i=0}^{k-1} C_i + \frac{1}{k} \left(\sum_{i=0}^{k-1} \Delta t_i C_i^{\frac{1}{2}} \right) + \delta_{r=0} \gamma^{-1} \right),$$

for some constant $C > 0$. First, we notice that

$$\begin{aligned}\eta_k^r &= (I_d - \gamma H)^{k-1} \eta_0^r + \sum_{i=0}^{k-1} (I_d - \gamma H)^{k-1-i} (\chi_i^r + \Delta t_i \psi_i^r) \\ \bar{\eta}_k^r &= \frac{1}{\gamma k} H^{-1} \left(I_d - (I_d - \gamma H)^k \right) \eta_0^r + \frac{1}{\gamma k} \sum_{i=0}^{k-1} \left(I_d - (I_d - \gamma H)^{k-i} \right) H^{-1} (\chi_i^r + \Delta t_i \psi_i^r),\end{aligned}$$

which implies that

$$\begin{aligned}\mathbb{E} \left[(\bar{\eta}_k^r)^\top S \bar{\eta}_k^r \right] &\leq \frac{3}{\gamma^2 k^2} (\eta_0^r)^\top \left(I_d - (I_d - \gamma H)^k \right)^\top H^{-\top} S H^{-1} \left(I_d - (I_d - \gamma H)^k \right) \eta_0^r \\ &\quad + \frac{3}{\gamma^2 k^2} \sum_{i=0}^{k-1} \mathbb{E} \left[\chi_i^\top \left(I_d - (I_d - \gamma H^\top)^{k-i} \right) H^{-\top} S H^{-1} \left(I_d - (I_d - \gamma H)^{k-i} \right) \chi_i \right] \\ &\quad + \frac{3}{\gamma^2 k^2} \sum_{0 \leq i, j \leq k-1} \mathbb{E} \left[\psi_i^\top \left(I_d - (I_d - \gamma H^\top)^{k-i} \right) H^{-\top} S H^{-1} \left(I_d - (I_d - \gamma H)^{k-j} \right) \psi_j \right].\end{aligned}$$

Let us define $I_{k,0}^r$, $I_{k,1}^r$ and $I_{k,2}^r$ as the first, second and third term, respectively, in the right-hand side of the latter inequality. One may notice that $I_{k,0}^r = 0$ if $r \geq 1$. Then concerning, $I_{k,0}^0$, we get

$$\begin{aligned}I_{k,0}^0 &= \frac{3}{2\gamma^2 k^2} \eta_0^\top \left(I_d - (I_d - \gamma H^\top)^{k-i} \right) \left(H^{-\top} + H^{-1} \right) \left(I_d - (I_d - \gamma H)^{k-i} \right) \eta_0 \\ &\leq \frac{C}{\gamma^2 k} \eta_0^\top \eta_0 \leq \frac{C}{\gamma^2 k},\end{aligned}$$

where we used (B.9) to obtain the last line. Then let us pass to $I_{k,1}^r$,

$$\begin{aligned}I_{k,1}^r &= \frac{3}{2\gamma^2 k^2} \sum_{i=0}^{k-1} \mathbb{E} \left[(\chi_i^r)^\top \left(I_d - (I_d - \gamma H^\top)^{k-i} \right) \left(H^{-\top} + H^{-1} \right) \left(I_d - (I_d - \gamma H)^{k-i} \right) \chi_i^r \right] \\ &= \frac{3}{2\gamma^2 k^2} \text{tr} \sum_{i=0}^{k-1} \left(I_d - (I_d - \gamma H)^{k-i} \right) \mathbb{E} [\chi_i^r \otimes \chi_i^r] \left(I_d - (I_d - \gamma H^\top)^{k-i} \right) \left(H^{-\top} + H^{-1} \right) \\ &\leq \frac{3\gamma^{r-1} R^{2r}}{2k^2} \text{tr} \sum_{i=0}^{k-1} C_i \left(I_d - (I_d - \gamma H)^{k-i} \right) S \left(I_d - (I_d - \gamma H^\top)^{k-i} \right) \left(H^{-\top} + H^{-1} \right) \\ &= \frac{3\gamma^{\max(r-1,0)} R^{2r}}{2k^2} \text{tr} \sum_{i=0}^{k-1} C_i \left(I_d - (I_d - \gamma H^\top)^{k-i} \right) \left(I_d - (I_d - \gamma H)^{k-i} \right) (2I_d + H H^{-\top} + H^{-1} H^\top) \\ &\leq \frac{C\gamma^{\max(r-1,0)} R^{2r}}{k^2} \text{tr} (I_d + H H^{-\top}) \sum_{i=0}^{k-1} C_i.\end{aligned}$$

Then, concerning $I_{k,2}^r$, using the triangular inequality, we get

$$\begin{aligned}I_{k,2}^r &\leq \frac{3}{2\gamma^2 k^2} \left(\sum_{i=0}^{k-1} \Delta t_i \mathbb{E} \left[(\psi_i^r)^\top \left(I_d - (I_d - \gamma H^\top)^{k-i} \right) \left(H^{-\top} + H^{-1} \right) \left(I_d - (I_d - \gamma H)^{k-i} \right) \psi_i^r \right]^{\frac{1}{2}} \right)^2 \\ &\leq \frac{C\gamma^{\max(r-1,0)} R^{2r}}{2k^2} \left(\sum_{i=0}^{k-1} \Delta t_i \left[C_i \text{tr} \left(I_d + H^{-\top} H \right) \right]^{\frac{1}{2}} \right)^2 \\ &= \frac{C\gamma^{\max(r-1,0)} R^{2r}}{2k^2} \text{tr} \left(I_d + H^{-\top} H \right) \left(\sum_{i=0}^{k-1} \Delta t_i C_i^{\frac{1}{2}} \right)^2,\end{aligned}$$

where we obtained the second line with similar arguments as in the calculus of the bound of $I_{k,1}^r$ above.

Third step: getting the desired bound. Using the triangular inequality on the norm induced by S , we obtain

$$\begin{aligned} \mathbb{E} \left[(\bar{\eta}_k)^\top S \bar{\eta}_k \right] &\leq \left(\sum_{r=0}^{k-1} \mathbb{E} \left[(\bar{\eta}_k^r)^\top S \bar{\eta}_k^r \right]^{\frac{1}{2}} \right)^2 \\ &\leq \frac{C}{\gamma k} + \frac{C}{k^2(1-\gamma^{\frac{1}{2}}R)} \text{tr}(I_d + HH^{-\top}) \left(\sum_{i=0}^{k-1} C_i + \left(\sum_{i=0}^{k-1} \Delta t_i C_i^{\frac{1}{2}} \right)^2 \right). \end{aligned}$$

Therefore, if $\sum_{k=0}^{\infty} \Delta t_k^2$ is finite, then C_k is uniformly bounded and we can conclude by taking $\varepsilon = \Delta t_0^{-2}$. If instead $\sum_{i=0}^{k-1} \Delta t_i^2 \leq a \ln(1+k)$, we obtain that $C_k \leq (1+k)^{a\varepsilon}$ and $\sum_{i=0}^{k-1} C_i$ is of order $k^{1+a\varepsilon}$ leading to the desired inequality up to changing ε into $a^{-1}\varepsilon$. \square

B Technical lemmas

B.1 Expansions of the temporal differences

Lemma B.1. *For $(x, \xi, \theta) \in \Omega \times \mathbb{R}^d \times \Theta$ and $0 < \Delta t < 1$, there exist $R(x, \xi, \theta)$ such that*

$$\begin{aligned} \delta_{\Delta t}(x, \xi, \theta) &= \mathcal{L}v(x) + R_0(x, \xi)^\top \theta + \Delta t^{\frac{1}{2}} R_1(x, \xi)^\top \theta + \Delta t R_2(\Delta t, x, \xi)^\top \theta \\ R_0(x, \xi)^\top \theta &= \frac{\sigma^2}{2} \left(\Delta_x v(x) - \xi^\top D_x^2 v(x) \xi \right), \\ R_1(x, \xi)^\top \theta &= \rho \sigma \nabla_x v(x) \cdot \xi - \frac{\sigma}{2} b(x, u(x))^\top D_x^2 v(x) \xi - \frac{\sigma^3}{6} d_x^3 v(x)(\xi, \xi, \xi), \end{aligned}$$

for some $R_2(\Delta t, x, \xi)$ such that, if ξ is a random variable normally distributed with mean zero and variance one, then for $p \geq 1$, $\mathbb{E}[|R_2(\Delta t, x, \xi)|^p]$ is bounded uniformly with respect to Δt and x .

Proof. The proof consists in defining $\varphi : [0, 1] \rightarrow \mathbb{R}$ by

$$\varphi(s) = e^{-s\rho\Delta t} v \left(x + s \left(b(x, u(x))\Delta t + \sigma\sqrt{\Delta t}\xi \right) \right),$$

and taking the development up to order four,

$$\varphi(1) = \varphi(0) + \varphi'(0) + \frac{\varphi''(0)}{2} + \frac{\varphi'''(0)}{6} + \int_0^1 \frac{(1-s)^3}{6} \varphi''''(s) ds.$$

Using $\tilde{b} \in \mathbb{R}^d$ defined by $\tilde{b} = b(x, u(x))\Delta t + \sigma\sqrt{\Delta t}\xi$, the latter derivatives of φ are given by

$$\begin{aligned} \varphi(0) &= v(x) \\ \varphi'(0) &= -\rho\Delta t v(x) + \nabla_x v(x) \cdot \tilde{b} \\ \varphi''(0) &= \rho^2 \Delta t^2 v(x) - 2\rho\Delta t \nabla_x v(x) \cdot \tilde{b} + d_x^2 v(x)(\tilde{b}, \tilde{b}) \\ \varphi'''(0) &= -\rho^3 \Delta t^3 v(x) + 3\rho^2 \Delta t^2 \nabla_x v(x) \cdot \tilde{b} - 3\rho\Delta t d_x^2 v(x)(\tilde{b}, \tilde{b}) + d_x^3 v(x)(\tilde{b}, \tilde{b}, \tilde{b}) \\ \varphi''''(s) &= e^{-s\rho\Delta t} \left[\rho^4 \Delta t^4 v - 4\rho^3 \Delta t^3 \nabla_x v \cdot \tilde{b} + 6\rho^2 \Delta t^2 d_x^2 v(\tilde{b}, \tilde{b}) - 4\rho\Delta t d_x^3 v(\tilde{b}, \tilde{b}, \tilde{b}) + d^4 v(\tilde{b}, \tilde{b}, \tilde{b}, \tilde{b}) \right]. \end{aligned}$$

We conclude by replacing all the equalities in this proof in (1.7). \square

Lemma B.2. *There exists $C > 0$ such that, for any $(x, \theta) \in \Omega \times \Theta$, $n \geq 1$, $0 < \Delta t < \frac{1}{n}$ and $\xi = (\xi_i)_{0 \leq i < n}$ independent normally distributed random variables with mean zero and variance one, we have*

$$\begin{aligned} |\mathbb{E} [|\delta_{\Delta t}^n(x, \xi, \theta)|^2] - \mathcal{L}v(x)^2| &\leq C (1 + |\theta|^2) (n^{-1} + n\Delta t), \\ |\mathbb{E} [\nabla_{\theta} |\delta_{\Delta t}^n(x, \xi, \theta)|^2] - \nabla_{\theta} \mathcal{L}v(x)^2| &\leq C (1 + |\theta|) (n^{-1} + n\Delta t). \end{aligned}$$

Proof. Taking $X_0 = x$ and $X_{t_{i+1}} = S(X_{t_i}, \xi_i)$ for $0 \leq i < n$, we obtain

$$(B.1) \quad \delta_{\Delta t}^n(x, \xi, \theta) = \frac{1}{n} \sum_{i=0}^{n-1} \delta_{\Delta t}(X_i, \xi_i, \theta).$$

Let us do the expansion of $\mathcal{L}v(X_{t_i})$ around x up to order two,

$$\mathcal{L}v(X_{t_i}) = \mathcal{L}v(x) + \nabla_x \mathcal{L}v(x) \cdot \tilde{b} + \int_0^1 (1-s) \tilde{b}^\top D_x^2 \mathcal{L}v(x + s\tilde{b}) \tilde{b} ds,$$

where $\tilde{b} = \sum_{j=0}^{i-1} (b(X_{t_j}, u(X_{t_j}))\Delta t + \sigma\sqrt{\Delta t}\xi_j)$. The latter equalities and Lemma B.1 imply

$$\begin{aligned} \delta^n(x, \xi, \theta) &= \mathcal{L}(x) + \frac{\sigma^2}{2n} \sum_{i=0}^{n-1} (\Delta_x v(X_{t_i}) - \xi_i^\top D^2 v(X_{t_i}) \xi_i) + \frac{1}{n\sqrt{\Delta t}} \sum_{i=0}^{n-1} [(n-1-i)\sigma \nabla_x \mathcal{L}v(X_{t_i}) \cdot \xi_i \\ &\quad + \rho\sigma \nabla_x v(X_{t_i}) \cdot \xi_i - \frac{\sigma}{2} b(X_{t_i}, u(X_{t_i}))^\top D^2 v(X_{t_i}) \xi_i - \frac{\sigma^3}{6} d_x^3 v(X_{t_i})(\xi_i, \xi_i, \xi_i)] + R_{\Delta t}^n(x, \xi, \theta), \end{aligned}$$

with $\mathbb{E} [|\delta_{\Delta t}^n(x, \xi, \theta)|^2] \leq C(1 + |\theta|^2)(n^{-1} + n\Delta t)^2$. We conclude by taking the expectation of the square in the latter equality and using the independence of $(\xi_i)_{0 \leq i < n}$. \square

B.2 Some lemmas used in the proof of Theorem 3.3

Lemma B.3. *The matrices S and A are respectively the symmetric and asymmetric part of H . Moreover, they satisfy*

$$(B.2) \quad S^2 \leq \text{tr}(S)S$$

$$(B.3) \quad A^\top A = -A^2 \leq \frac{2}{\rho\sigma^2} \left\| b + \frac{\sigma^2}{2} \nabla_x \ln(m) \right\|_\infty^2 S^2$$

$$(B.4) \quad (SA - AS) \leq 2\sqrt{\frac{2}{\rho\sigma^2}} \left\| b + \frac{\sigma^2}{2} \nabla_x \ln(m) \right\|_\infty S^2,$$

$$(B.5) \quad \mathbb{E} [H(X)H(X)^\top] \leq \rho^{-1} \|\bar{\mathcal{L}}\varphi(X)\|_\infty^2 S.$$

Proof. *First step: proving that S and A are respectively the symmetric and asymmetric part of H .* Take $\theta \in \Theta$, we get:

$$\begin{aligned} \theta^\top H\theta &= \theta^\top \mathbb{E} [\varphi(X) \bar{\mathcal{L}}\varphi(X)^\top] \theta \\ &= \mathbb{E} [v(X, \theta) \bar{\mathcal{L}}v(X, \theta)] \\ &= \int_\Omega \left(\rho v - \frac{\sigma^2}{2} \Delta_x v + b(x) \cdot \nabla_x v \right) v(x) m(x) dx \\ &= \rho \mathbb{E} [v(X)^2] + \frac{\sigma^2}{2} \mathbb{E} [|\nabla_x v(X)|^2], \end{aligned}$$

where the last line is obtained using (2.2) and the following integration by parts,

$$\begin{aligned}
\int_{\Omega} \nabla_x v \cdot b(x) v(x) m(x) dx &= \int_{\Omega} \frac{1}{2} \nabla_x (v^2) \cdot b(x) m(x) dx \\
&= -\frac{1}{2} \int_{\Omega} \operatorname{div}(b(x) m(x)) v^2(x) dx, \\
-\int_{\Omega} \Delta_x v(x) v(x) m(x) dx &= \int_{\Omega} |\nabla_x v|^2 m(x) dx + \int_{\Omega} \frac{1}{2} \nabla_x (v^2) \cdot \nabla_x m(x) dx \\
&= \int_{\Omega} |\nabla_x v|^2 m(x) dx - \frac{1}{2} \int_{\Omega} \Delta_x m(x) v^2(x) dx.
\end{aligned}$$

This implies that S is the symmetric part of H . Then it is straightforward that the asymmetric part of H is equal to A .

Second step: proving the four inequalities. The first inequality (B.2) is straightforward, it only relies on the fact that S is symmetric and positive. The fourth inequality (B.5) is straightforward using the definitions of $H(X)$ and S . The third inequality (B.4) is a consequence of (B.3). Therefore, there is only (B.3) left to prove. Let us take $\lambda \in \mathbb{C}$ a complex eigenvalue of H , and θ an associated normalized eigenvector, it satisfies $\bar{\theta}^\top S \theta = \Re(\lambda)$ and $\bar{\theta}^\top A \theta = i \Im(\lambda)$. Therefore, we get

$$\begin{aligned}
|\Im(\lambda)| &= \left| \bar{\theta}^\top A \theta \right| \\
&= \left| \mathbb{E} \left[\bar{v}(X, \theta) (b(X) + \nabla_x \ln m(X))^\top \nabla_x v(X, \theta) \right] \right| \\
&\leq \|b + \nabla_x \ln(m)\|_{\infty} \mathbb{E} [|v(X, \theta)|^2]^{\frac{1}{2}} \mathbb{E} [|\nabla_x v(X, \theta)|^2]^{\frac{1}{2}} \\
&\leq \sqrt{\frac{2}{\rho \sigma^2}} \|b + \nabla_x \ln(m)\|_{\infty} \bar{\theta}^\top S \theta.
\end{aligned}$$

This concludes the proof. \square

Lemma B.4. *For $\gamma \leq R^{-2}$, the following two inequalities hold for any $k \geq 0$,*

$$(B.6) \quad (I_d - \gamma H^\top)(I_d - \gamma H) \leq I_d - \gamma S$$

$$(B.7) \quad \left(I_d - (I - \gamma H^\top)^k \right) \left(I_d - (I - \gamma H)^k \right) \leq \gamma^2 k^2 H^\top H,$$

$$(B.8) \quad \left(I_d - (I - \gamma H^\top)^k \right) \left(I_d - (I - \gamma H)^k \right) \leq 4 \left(1 + \frac{2}{\rho \sigma^2} \|b + \nabla_x \ln(m)\|_{\infty}^2 \right) I_d,$$

(B.9)

$$\left(I_d - (I - \gamma H^\top)^k \right) \left(H^{-1} + H^{-\top} \right) \left(I_d - (I - \gamma H)^k \right) \leq 2\gamma k \left(1 + \sqrt{\frac{2}{\rho \sigma^2}} \|b + \nabla_x \ln(m)\|_{\infty} \right) I_d.$$

The latter lemma would be straight forward if H were symmetric. Conversely, it does not hold if we only assume the eigenvalues of H to be bounded and with positive real part. In fact, we need some bound on the imaginary part of the spectrum of H , depending on its real part.

Proof. One may notice that (B.9) is a straightforward consequence of (B.7) and (B.8). Then, concerning (B.6), it is sufficient to write $(I_d - \gamma H^\top)(I_d - \gamma H) = I_d - 2\gamma S + \gamma^2 (S^2 + SA - AS - A^2)$, and use the definition of R , (B.2), (B.3) and (B.4). Therefore, it only remains to prove (B.7) and (B.8).

First step: proving (B.7). Let us proceed by induction, the case $k = 0$ is straightforward. Let us denote $y_k = \left(I_d - (I - \gamma H)^k \right)$ and assume that the inequality holds for k . One may

notice that for $\theta \in \mathbb{R}^d$, using (B.6), we obtain

$$\begin{aligned} \theta^\top y_k^\top (I_d - \gamma H)^\top H \theta &\leq \left(\theta^\top y_k^\top (I_d - \gamma H)^\top (I_d - \gamma H) y_k \theta \right)^{\frac{1}{2}} \left(\theta^\top H^\top H \theta \right)^{\frac{1}{2}} \\ &\leq \gamma k \theta^\top H^\top H \theta, \end{aligned}$$

which implies $y_k^\top (I_d - \gamma H)^\top H + H^\top (I_d - \gamma H) y_k \leq 2\gamma k H^\top H$. Using the latter inequality, the relation $y_{k+1} = (I_d - \gamma H) y_k + \gamma H$, and (B.6) again, we get

$$\begin{aligned} y_{k+1}^\top y_{k+1} &= y_k^\top (I_d - \gamma H)^\top (I_d - \gamma H) y_k + \gamma y_k^\top (I_d - \gamma H)^\top H + \gamma H^\top (I_d - \gamma H) y_k + \gamma^2 H^\top H \\ &\leq \gamma^2 k^2 H^\top H + 2\gamma^2 k H^\top H + \gamma^2 H^\top H = \gamma^2 (k+1)^2 H^\top H. \end{aligned}$$

This concludes the induction.

Second step: proving (B.8). In this step, we will only work with the complex eigenvalues of H : let $\lambda \in \mathbb{C}$ be one of them, we get

$$\begin{aligned} \left| 1 - (1 - \gamma\lambda)^{k+1} \right| &= \left| (1 - \gamma\lambda) \left(1 - (1 - \gamma\lambda)^k \right) + \gamma\lambda \right| \\ &\leq \left(|1 - \gamma\lambda| \left| 1 - (1 - \gamma\lambda)^k \right| + \gamma|\lambda| \right). \end{aligned}$$

This implies

$$\begin{aligned} \left| 1 - (1 - \gamma\lambda)^k \right| &\leq \gamma|\lambda| \sum_{j=0}^{k-1} |1 - \gamma\lambda|^j \\ &\leq \frac{\gamma|\lambda|}{1 - |1 - \gamma\lambda|} \\ &\leq \frac{\gamma|\lambda|}{1 - (1 - \gamma\Re(\lambda))^{\frac{1}{2}}} \quad \text{using (B.6),} \\ &\leq \frac{\gamma|\lambda|}{1 - (1 - \frac{\gamma}{2}\Re(\lambda))} \quad \text{because } \gamma\Re(\lambda) \leq 1, \\ &\leq 2\sqrt{1 + \frac{\Im(\lambda)^2}{\Re(\lambda)^2}} \\ &\leq 2 \left(1 + \frac{2}{\rho\sigma^2} \|b + \nabla_x \ln(m)\|_\infty^2 \right)^{\frac{1}{2}}, \end{aligned}$$

where the last inequality comes from a similar argument as in the proof of (B.3). This concludes the proof. \square

Lemma B.5. *Assume H3. There exists $C > 0$ such that the two following inequalities hold for any $k \geq 0$,*

$$\begin{aligned} \mathbb{E} [\varphi(X_k) \otimes \varphi(X_k)] &\leq CS, \\ (\mathbb{E} [H(X_k)] - H)(\mathbb{E} [H(X_k)] - H)^\top &\leq C\Delta t_k^2 S. \end{aligned}$$

Proof. We recall that the set of admissible functions v is finitely dimensional, therefore the C^4 -norm and the $H^1(m)$ -norm are equivalent and there exists $C > 0$ such that $\|v(\cdot, \theta)\|_{C^4}^2 \leq C\|v(\cdot, \theta)\|_{H^1(m)}^2$. For $\theta \in \Theta$ and $k \geq 0$, this implies

$$\begin{aligned} \theta^\top \mathbb{E} [\varphi(X_k) \otimes \varphi(X_k)] \theta &= C\mathbb{E} [v(X_k, \theta)^2] \\ &\leq C\mathbb{E} [v(X, \theta)^2] + C\Delta t_k \|v(\cdot, \theta)\|_{C^4}^2 \\ &\leq C(1 + \Delta t_k) \|v(\cdot, \theta)\|_{H^1(m)}^2, \end{aligned}$$

where the second line is obtained from Theorem 3.1. Here, C is a constant that can change from line to line. The first inequality is then obtained by recalling that $\|v(\cdot, \theta)\|_{H^1(m)}^2 \leq (\rho^{-1} + 2\sigma^{-2})\theta^\top S\theta$.

Concerning the second inequality, we get

$$\begin{aligned} |(\mathbb{E}[H(X_k)] - H)\theta|^2 &= |\mathbb{E}[\varphi(X_k)\mathcal{L}v(X_k, \theta) - \varphi(X)\mathcal{L}v(X, \theta)]|^2 \\ &\leq C(\Delta t_k \|v(\cdot, \theta)\|_{C^6})^2 \\ &\leq C\Delta t_k^2 \|v(\cdot, \theta)\|_{H^1(m)}^2, \end{aligned}$$

where the second line is obtained from Theorem 3.1, and the third line from the fact that the C^6 -norm is equivalent to the $H^1(m)$ on the finite dimensional space of functions v . We conclude the same way as we did for the first inequality. \square

B.3 Calculus of variances and covariances

Lemma B.6. *Let $(x, \theta) \in \Omega \times \Theta$ and ξ a Gaussian vector with mean zero and variance one, the following equalities hold*

$$(B.10) \quad \text{Var}(\xi \cdot \nabla_x v(x)) = |\nabla_x v(x)|^2,$$

$$(B.11) \quad \text{Var}(\xi^\top D^2 v(x) \xi - \Delta_x v(x)) = \text{tr}(D_x^2 v(x)^2).$$

Proof. The first equality is straightforward. Since $D^2 v(x)$ is symmetric, there exists P an orthogonal matrix and D a diagonal matrix such that $D^2 v(x) = P^\top D P$. The couples (X, ξ) and $(X, P^\top \xi)$ have the same law and ξ is independent of X and D , this implies

$$\begin{aligned} \text{Var}(\xi^\top D^2 v(x) \xi - \Delta_x v(x)) &= \mathbb{E} \left[\left(\xi^\top D^2 v(x) \xi - \Delta_x v(x) \right)^2 \right] \\ &= \mathbb{E} \left[\left((P^\top \xi)^\top D^2 v(x) P^\top \xi - \Delta_x v(x) \right)^2 \right] \\ &= \mathbb{E} \left[\left(\xi^\top D \xi - \Delta_x v(x) \right)^2 \right] \\ &= \mathbb{E} \left[\sum_{i=1}^d D_i^2 (\xi_i^2 - 1)^2 \right] = 2 \sum_{i=1}^d D_i^2 = 2 \text{tr}(D_x^2 v(x)^2). \end{aligned}$$

This concludes the proof. \square