

Robust SDE-Based Variational Formulations for Solving Linear PDEs via Deep Learning

Lorenz Richter^{*123} Julius Berner^{*4}

Abstract

The combination of Monte Carlo methods and deep learning has recently led to efficient algorithms for solving partial differential equations (PDEs) in high dimensions. Related learning problems are often stated as variational formulations based on associated stochastic differential equations (SDEs), which allow the minimization of corresponding losses using gradient-based optimization methods. In respective numerical implementations it is therefore crucial to rely on adequate gradient estimators that exhibit low variance in order to reach convergence accurately and swiftly. In this article, we rigorously investigate corresponding numerical aspects that appear in the context of linear Kolmogorov PDEs. In particular, we systematically compare existing deep learning approaches and provide theoretical explanations for their performances. Subsequently, we suggest novel methods that can be shown to be more robust both theoretically and numerically, leading to substantial performance improvements.

1. Introduction

In this paper we suggest novel methods for solving high-dimensional linear Kolmogorov PDEs. In particular, those methods allow for more efficient approximations by reducing the variance of loss estimators, as illustrated by numerical experiments like the one displayed in Figure 1.

High-dimensional PDEs are ubiquitous in applications appearing in economics, science, and engineering, however,

^{*}Equal contribution (the author order was determined by `numpy.random.rand(1)`)¹dida Datenschmiede GmbH, Germany ²Zuse Institute Berlin, Germany ³Institute of Mathematics, Freie Universität Berlin, Germany ⁴Faculty of Mathematics, University of Vienna, Austria. Correspondence to: Julius Berner <julius.berner@univie.ac.at>, Lorenz Richter <lorenz.richter@fu-berlin.de>.




Figure 1. Standard deviation and performance (in terms of the mean squared error (MSE) on evaluation data) of different losses as a function of the batch size after training a neural network for $30k$ steps, where due to memory constraints different losses allow for different maximal batch size values.

their numerical treatment poses formidable challenges since traditional grid-based methods suffer from the curse of dimensionality. Recently, the combination of Monte Carlo methods and deep learning has led to feasible algorithms (E et al., 2017; Sirignano & Spiliopoulos, 2018; Han et al., 2018; Raissi et al., 2019; Pfau et al., 2020; Berner et al., 2021), which, at least in some settings, are able to beat this curse (Jentzen et al., 2018; Reisinger & Zhang, 2019; Berner et al., 2020b). The general idea is to rely on stochastic representations of the PDE which may be interpreted as variational formulations that can subsequently be minimized by iterative gradient decent methods (Nüsken & Richter, 2021a). In respective algorithms it is therefore crucial to rely on adequate gradient estimators that exhibit low variance and eventually lead to robust numerical optimization routines, where robustness refers to the ability of coping with stochasticity in the optimization process.

In this article we focus on linear Kolmogorov PDEs for which variational formulations may be based on the celebrated Feynman-Kac formula (Beck et al., 2021). This method is handy from an implementational point of view, however, we can demonstrate potential non-robustness issues that might deteriorate the numerical performance. To overcome these issues, we suggest alternative approaches resting on backward stochastic differential equations which lead to variance-reduced estimators and robust algorithms. Our contributions can be summarized as follows:

- We provide a systematic numerical comparison of loss

functions for solving linear Kolmogorov PDEs via deep learning.

- We can show both theoretically and numerically why certain methods are non-robust and might therefore incur numerical instabilities.
- We provide new loss functions and optimization methods which are provably robust and ultimately yield significant performance improvements in relevant high-dimensional numerical experiments.

1.1. Related Work

Solving linear Kolmogorov PDEs via the Feynman-Kac formula by means of deep learning has been suggested in Beck et al. (2021) and further analyzed in Berner et al. (2020a). For an analysis of certain kinds of elliptic linear PDEs we refer to Grohs & Herrmann (2020). Variational formulations of nonlinear PDEs based on backward stochastic differential equations (BSDEs) have been pioneered in E et al. (2017) and extensions that aim for approximations on entire domains have, for instance, been suggested in Nüsken & Richter (2021a). We refer to Richter (2021) for a comprehensive introduction and further perspectives and to Beck et al. (2019) for the treatment of fully nonlinear equations. An alternative method for the approximation of high-dimensional PDEs is termed *physics-informed neural networks* (PINNs) (Raissi et al., 2019; Sirignano & Spiliopoulos, 2018), which relies on iterative residual minimizations. Note that in contrast to most SDE-based methods it relies on the explicit computation of Hessians and involves additional loss terms for the boundary values, which need to be tuned carefully.

In terms of variance reduction of Monte Carlo estimators we refer to Vidales et al. (2018) for a control variate attempt via deep learning. Robustness and variance analyses of certain loss functions related to PDE problems have been conducted in Nüsken & Richter (2021b). The related observation that including a certain Itô integral term into the objective function can lead to variance-reduced estimators has been made in Zhou et al. (2021). Both works, however, consider different settings and focus only on Hamilton–Jacobi–Bellman (HJB) equations. Furthermore, they do not provide a detailed analysis of the interplay between robustness properties, performance, and computational feasibility, which is crucial for practical applications.

1.2. Notation

We assume that our random variables are defined on a suitable underlying filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_t, \mathbb{P})$ satisfying the usual conditions (Klenke, 2013). We denote the expectation of a random variable $A: \Omega \rightarrow \mathbb{R}$ by $\mathbb{E}[A] := \int A \, d\mathbb{P}$. Given another random variable $B: \Omega \rightarrow \mathbb{R}^k$ with $k \in \mathbb{N}$, we also consider the condi-

tional expectation $\mathbb{E}[A|B]$ of A given B and the conditional expectation $b \mapsto \mathbb{E}[A|B = b]$ of A given $B = b$, see, e.g., Ash & Doléans-Dade (2000) for the definitions. We further define the variance of A by $\mathbb{V}[A] := \mathbb{E}[(A - \mathbb{E}[A])^2]$ and the conditional variance of A given B by $\mathbb{V}[A|B] := \mathbb{E}[(A - \mathbb{E}[A|B])^2|B]$. Throughout the article, W is a standard d -dimensional Brownian motion which is adapted to the underlying filtration and we consider stochastic integrals w.r.t. W as, e.g., defined in Gall (2016). Further remarks on the notation can be found in Appendix A.

2. SDE-Based Variational Formulations of Linear PDEs

In this paper we consider linear *Kolmogorov partial differential equations*¹ of the type

$$(\partial_t + L)V(x, t) = 0, \quad (x, t) \in \mathbb{R}^d \times [0, T], \quad (1a)$$

$$V(x, T) = g(x), \quad x \in \mathbb{R}^d, \quad (1b)$$

where $g \in C(\mathbb{R}^d, \mathbb{R})$ is a given terminal condition and

$$L := \frac{1}{2} \sum_{i,j=1}^d (\sigma \sigma^\top)_{ij}(x, t) \partial_{x_i} \partial_{x_j} + \sum_{i=1}^d b_i(x, t) \partial_{x_i} \quad (2)$$

is a differential operator based on given coefficient functions $\sigma \in C(\mathbb{R}^d \times [0, T], \mathbb{R}^{d \times d})$ and $b \in C(\mathbb{R}^d \times [0, T], \mathbb{R}^d)$. For simplicity, let us assume the existence of a unique strong solution $V \in C^{2,1}(\mathbb{R}^d \times [0, T], \mathbb{R})$ to the PDE in (1) and refer to Assumption 1 in Appendix B for further technical details. At the same time, let us note that most of our results also hold for the more general concept of viscosity solutions that are continuous but not necessarily differentiable and require weaker conditions on the functions b, σ , and g , see Hairer et al. (2015). Furthermore, we can also consider more general linear PDEs as well as elliptic and parabolic problems on bounded domains, as elaborated on in Appendix C.

This covers a broad spectrum of PDEs, for which accurate and reliable solvers are of great importance to practitioners. For instance, such PDEs frequently appear in physics for the modelling of heat flow and diffusion processes (Widder, 1976; Pascucci, 2005). Moreover, the PDE in (1) includes the Black-Scholes equation and extensions thereof, used for pricing financial derivative instruments (Black & Scholes, 1973; Pironneau & Achdou, 2009; Ekström & Tysk, 2010). Finally, we want to emphasize appearances of such PDEs in the field of machine learning, e.g., in the context of reinforcement learning (Theodorou et al., 2010; Kiumarsi et al., 2017) and diffusion-based generative modeling (Huang et al., 2021).

¹PDEs of this type are often also referred to as *Kolmogorov (backward) equations* and their adjoints are known as *Fokker-Planck equations* or *Kolmogorov forward equations*.

In order to design machine learning algorithms aiming to approximate solutions to (1), we consider loss functionals

$$\mathcal{L} : \mathcal{U} \rightarrow \mathbb{R}_{\geq 0}, \quad (3)$$

which shall be minimal if and only if $u \in \mathcal{U}$ fulfills (1), thereby offering a variational formulation that allows for iterative minimization strategies. Here, \mathcal{U} is an appropriate function class holding sufficient approximation capacity and containing, for instance, a class of certain neural networks. Moreover, we implicitly assume that \mathcal{U} is chosen such that $V \in \mathcal{U}$ and, for our theoretical results, we require the functions in \mathcal{U} to be sufficiently smooth.

A first obvious choice for the loss could be the squared difference between the approximating function u and the solution V , i.e.

$$\mathcal{L}_{\text{Eval}}(u) := \mathbb{E} \left[(V(\xi, \tau) - u(\xi, \tau))^2 \right], \quad (4)$$

where ξ and τ are suitable² random variables with values in \mathbb{R}^d and $[0, T]$, respectively. At first glance, this loss seems intractable since the solution V is just the quantity we are after and is therefore not known. However, let us recall the *Feynman-Kac formula*, which connects the deterministic function V to a stochastic process by stating that almost surely it holds that

$$V(\xi, \tau) = \mathbb{E} [g(X_T) | (\xi, \tau)], \quad (5)$$

where X is the solution to the SDE

$$dX_s = b(X_s, s) ds + \sigma(X_s, s) dW_s, \quad X_\tau = \xi, \quad (6)$$

see Appendix C for details. This allows us to rewrite (4) as

$$\mathcal{L}_{\text{Eval}}(u) = \mathbb{E} \left[\left(\mathbb{E} [g(X_T) | (\xi, \tau)] - u(\xi, \tau) \right)^2 \right] \quad (7a)$$

$$= \mathbb{E} \left[\mathbb{E} [\Delta_u | (\xi, \tau)]^2 \right] \quad (7b)$$

$$= \mathbb{E} \left[\mathbb{E} [\Delta_u - S_u | (\xi, \tau)]^2 \right], \quad (7c)$$

which now relies on the (random) quantities³

$$\Delta_u := g(X_T) - u(\xi, \tau), \quad (8a)$$

$$S_u := \int_\tau^T (\sigma^\top \nabla_x) u(X_s, s) \cdot dW_s. \quad (8b)$$

²In practice, one often chooses random variables that are uniformly distributed on a given hypercube. For our theoretical results, we assume that ξ and τ are \mathcal{F}_0 -measurable and exhibit sufficiently many bounded moments. This ensures that ξ and τ are sufficiently concentrated and independent of the Brownian motion W .

³For the sake of readability, we omit several implicit dependencies. Specifically, the stochastic process X and the quantities Δ_u and S_u depend on the random variable (ξ, τ) , defining the initial time and value of the SDE, i.e., $X_\tau = \xi$. Also note that the stochastic integral S_u and the SDE in (6) are driven by the same Brownian motion W .

The last step (7c) follows from the fact that, under mild regularity assumptions, the stochastic integral S_u is a martingale which has zero expectation given (ξ, τ) . Adding quantities with a known expectation is a common trick for variance reduction of corresponding estimators known as *control variates*, see, e.g., Section 4.4.2 in Robert & Casella (2004). The benefit of particularly choosing S_u becomes clear with the following Lemma, which states that for $u = V$ the corresponding control variate actually yields a zero-variance estimator.

Lemma 2.1 (Optimal control variate). *Let V be a solution to the PDE in (1). For Δ_V and S_V as defined in (8a) and (8b) it almost surely holds that*

$$\Delta_V = S_V, \quad (9)$$

which implies that $\mathbb{V} [\Delta_V - S_V | (\xi, \tau)] = 0$.

Proof. The proof is an application of Itô's lemma. Together with the observation that X is an Itô process given by the SDE in (6), it implies that almost surely it holds that

$$V(X_T, T) - V(\xi, \tau) = \int_\tau^T (\partial_t + L)V(X_s, s) ds + S_V,$$

assuming that $V \in C^{2,1}(\mathbb{R}^d \times [0, T], \mathbb{R})$, see, e.g., Theorem 3.3.6 in Karatzas & Shreve (1998). Using the fact that V solves the PDE in (1), this implies the statement. \square

While, in principle, the new representation in (7) makes the loss $\mathcal{L}_{\text{Eval}}$ computationally tractable, an immediate disadvantage are the two expectations, which, for a numerical implementation, might be costly as they would both need to rely on Monte Carlo approximations.

It turns out, however, that the conditional expectations in (7b) and (7c) are in fact not needed and we can further define the two losses

$$\mathcal{L}_{\text{FK}}(u) := \mathbb{E} [\Delta_u^2] \quad (10)$$

and

$$\mathcal{L}_{\text{BSDE}}(u) := \mathbb{E} [(\Delta_u - S_u)^2]. \quad (11)$$

Note that now the expectations in (10) and (11) are to be understood with respect to random initial time τ , random initial value ξ , and the randomness of the Brownian motion W , which defines the evolution of the stochastic process X on the (random) interval $[\tau, T]$.

In the remainder of this article we will investigate the above losses with respect to certain robustness properties that will turn out to imply different statistical properties and consequently lead to different optimization performances of corresponding algorithms. Let us start with the following proposition which relates the three losses to one another and

therefore guarantees that \mathcal{L}_{FK} and $\mathcal{L}_{\text{BSDE}}$ indeed constitute meaningful objectives for approximating the solution V to the PDE in (1).

Proposition 2.2 (Minima of \mathcal{L}_{FK} and $\mathcal{L}_{\text{BSDE}}$). *For every measurable $u: \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}$ it holds that*

$$\mathcal{L}_{\text{Eval}}(u) = \mathcal{L}_{\text{FK}}(u) - \mathbb{V}[S_V] \quad (12a)$$

$$= \mathcal{L}_{\text{BSDE}}(u) - \mathbb{V}[S_{V-u}]. \quad (12b)$$

This implies that the solution V to the PDE in (1) is the unique⁴ minimizer of \mathcal{L}_{FK} and $\mathcal{L}_{\text{BSDE}}$. The minima satisfy

$$(i) \min_u \mathcal{L}_{\text{FK}}(u) = \mathcal{L}_{\text{FK}}(V) = \mathbb{E}[\Delta_V^2] = \mathbb{V}[S_V],$$

$$(ii) \min_u \mathcal{L}_{\text{BSDE}}(u) = \mathcal{L}_{\text{BSDE}}(V) = 0.$$

Proof. See Appendix B. \square

Remark 2.3 (Origin of the losses). As already hinted at, the names of the losses originate from the Feynman-Kac (FK) formula and backward stochastic differential equations (BSDE), respectively. The former, as stated in (5) and detailed in Appendix C, implies that the solution of the PDE in (1) can be expressed as a conditional expectation, i.e.,

$$V(x, t) = \mathbb{E}[g(X_T) | X_t = x] \quad (13a)$$

$$= \mathbb{E}[g(X_T) | (\xi, \tau) = (x, t)]. \quad (13b)$$

This can then be combined with the fact that the conditional expectation is the minimizer of a corresponding L^2 -optimization problem and one readily recovers \mathcal{L}_{FK} as specified in (10), see Beck et al. (2021). The latter loss, $\mathcal{L}_{\text{BSDE}}$, originates from the BSDE

$$dY_s = Z_s \cdot dW_s, \quad Y_T = g(X_T), \quad (14)$$

which can essentially be derived from Itô's lemma and which involves the backward processes $Y_s = V(X_s, s)$ and $Z_s = (\sigma^\top \nabla_x V)(X_s, s)$, see Appendix D for further details. The name of the loss $\mathcal{L}_{\text{Eval}}$ refers to the fact that it is mostly used in order to evaluate the solution candidate u , using either a known solution V or the reformulation in (7). This is also how we evaluate the mean squared error (MSE) for our methods, see Appendix E.

2.1. Estimator Versions of Losses and Robustness Issues

In practice, the expectations in the losses (4), (7), (10), or (11) can usually not be computed analytically and we need to resort to Monte Carlo approximations $\hat{\mathcal{L}}^{(K)}$ based on a given number $K \in \mathbb{N}$ of independent samples

$$(\xi^{(k)}, \tau^{(k)}, W^{(k)}) \sim (\xi, \tau, W), \quad k = 1, \dots, K.$$

⁴Up to null sets w.r.t. to the image measure of (ξ, τ) .

In particular, this yields the estimators

$$\hat{\mathcal{L}}_{\text{FK}}^{(K)}(u) = \frac{1}{K} \sum_{k=1}^K \left(\hat{\Delta}_u^{(k)} \right)^2, \quad (15a)$$

$$\hat{\mathcal{L}}_{\text{BSDE}}^{(K)}(u) = \frac{1}{K} \sum_{k=1}^K \left(\hat{\Delta}_u^{(k)} - \hat{S}_u^{(k)} \right)^2, \quad (15b)$$

where $\hat{\Delta}_u^{(k)}$ and $\hat{S}_u^{(k)}$ are discretized versions of (8a) and (8b), evaluated at the k -th sample. More precisely, we define

$$\hat{\Delta}_u^{(k)} := g(\hat{X}_{J^{(k)}+1}^{(k)}) - u(\xi^{(k)}, \tau^{(k)}), \quad (16a)$$

$$\hat{S}_u^{(k)} := \sum_{j=1}^{J^{(k)}} s_j^{(k)}, \quad (16b)$$

where

$$s_j^{(k)} := (\sigma^\top \nabla_x u)(\hat{X}_j^{(k)}, t_j^{(k)}) \cdot (W_{t_{j+1}^{(k)}} - W_{t_j^{(k)}}). \quad (17)$$

Here, $(\hat{X}_j^{(k)})_{j=1}^{J^{(k)}+1}$ denotes a discretization of the solution to the SDE in (6) (driven by the Brownian motion $W^{(k)}$) on a time grid $\tau^{(k)} = t_1^{(k)} < \dots < t_{J^{(k)}+1}^{(k)} = T$ with initial condition $\hat{X}_1^{(k)} = \xi^{(k)}$, see Appendix E for details.

In this section we will investigate different statistical properties of the above estimators. It will turn out that even though the loss \mathcal{L}_{FK} is prominent for solving linear PDEs in practice (see, e.g., Beck et al. (2021); Berner et al. (2020a)), it might have some disadvantages from a numerical point of view. In fact, Lemma 2.1 and Proposition 2.2 already show that, while $\mathcal{L}_{\text{BSDE}}$ can be minimized to zero expectation, \mathcal{L}_{FK} always exhibits the additional term

$$\mathbb{E}[\Delta_V^2] = \mathbb{V}[\Delta_V] = \mathbb{V}[S_V] = \mathbb{E}[S_V^2], \quad (18)$$

which is a potential source of additional noise. By the Itô isometry, this can also be written as

$$\mathbb{E}[S_V^2] = \mathbb{E} \left[\left(\int_\tau^T \|\sigma^\top \nabla_x V(X_s, s)\|^2 ds \right) \right]. \quad (19)$$

The following proposition shows what this implies for the estimator versions of the losses at the optimum.

Proposition 2.4 (Variance of losses). *For the estimator versions of the losses defined in (10) and (11) it holds that*

$$\mathbb{E} \left[\hat{\mathcal{L}}_{\text{FK}}^{(K)}(V) \right] = \mathbb{V}[S_V], \quad \mathbb{V} \left[\hat{\mathcal{L}}_{\text{FK}}^{(K)}(V) \right] = \frac{1}{K} \mathbb{V}[S_V^2]$$

and

$$\mathbb{E} \left[\hat{\mathcal{L}}_{\text{BSDE}}^{(K)}(V) \right] = 0, \quad \mathbb{V} \left[\hat{\mathcal{L}}_{\text{BSDE}}^{(K)}(V) \right] = 0.$$




Figure 2. Estimator losses and their standard deviations as a function of the gradient steps when solving the HJB equation from Section 3.3 with batch size $K = 128$. One observes that the statement of Proposition 2.4 already holds in an approximative sense before converging to the solution V .

Proof. See Appendix B. \square

Proposition 2.4 shows that neither the expectation nor the variance of $\widehat{\mathcal{L}}_{\text{FK}}^{(K)}$ can be zero at the solution $u = V$ unless S_V or S_V^2 are deterministic, which is usually not the case in practice as demonstrated by the numerical examples in Section 3 and Figure 2. While the non-zero expectation of the estimator is typically not a problem for the optimization routine, the non-vanishing variance might cause troubles when converging to the global minimum.

Since the considered losses are usually minimized via variants of gradient decent, our further analysis shall therefore focus on the variance of corresponding gradient estimators. For this let us first define an appropriate notion of derivative.

Definition 2.5 (Gateaux derivative). We say that $\mathcal{L}: \mathcal{U} \rightarrow \mathbb{R}_{>0}$ is Gateaux differentiable at $u \in \mathcal{U}$ if for all $\phi \in \mathcal{U}$ the mapping

$$\varepsilon \mapsto \mathcal{L}(u + \varepsilon\phi) \quad (20)$$

is differentiable at $\varepsilon = 0$. The Gateaux derivative of \mathcal{L} in direction ϕ is then defined as

$$\frac{\delta}{\delta u} \mathcal{L}(u; \phi) := \left. \frac{d}{d\varepsilon} \right|_{\varepsilon=0} \mathcal{L}(u + \varepsilon\phi). \quad (21)$$

Inspired by an analysis in Nüsken & Richter (2021b) let us first investigate the variance of derivatives at the solution $u = V$, where intuitively we might want to favor gradient estimators whose variances vanish at the global minimizer of the loss. The following proposition shows that this property is indeed only the case for $\widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}$.

Proposition 2.6 (Gradient variances). For the estimator versions of the losses defined in (10) and (11) and for all $\phi \in \mathcal{U}$ it holds that

$$(i) \quad \mathbb{V} \left[\left. \frac{\delta}{\delta u} \right|_{u=V} \widehat{\mathcal{L}}_{\text{FK}}^{(K)}(u; \phi) \right] = \frac{4}{K} \mathbb{V} [S_V \phi(\xi, \tau)],$$

$$(ii) \quad \mathbb{V} \left[\left. \frac{\delta}{\delta u} \right|_{u=V} \widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}(u; \phi) \right] = 0.$$

Proof. See Appendix B. \square




Figure 3. Gradient estimator standard deviations (maximum over directions $\phi = \partial_{\theta_i} \Phi_{\theta}$, where Φ_{θ} is a neural network with parameters θ , see Appendix E) and the MSE (on evaluation data) as a function of the gradient steps when solving the HJB equation from Section 3.3 with batch size $K = 128$. In line with Proposition 2.7 the standard deviation of the gradient of $\widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}$ decreases when the solution is sufficiently well approximated.

Note that even though the gradient estimator $\widehat{\mathcal{L}}_{\text{FK}}^{(K)}$ does not have vanishing variance at the solution, a large batch size K can counteract this effect. We will study corresponding dependencies between variances and batch sizes in our numerical experiments in Section 3.

While Proposition 2.6 only makes a statement on the gradient variances at the solution $u = V$, the following proposition shall indicate that, at least close to the solution, one might still expect a small variance of the gradients of $\widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}$, see also Figure 3.

Proposition 2.7 (Stability of $\mathcal{L}_{\text{BSDE}}$ close to the solution). Assume it holds almost surely that

$$|u(\xi, \tau) - V(\xi, \tau)| \leq \varepsilon \quad \text{and} \quad |\phi(\xi, \tau)| \leq \kappa \quad (22)$$

and that there exists $\gamma \in \mathbb{R}_{>0}$ such that for every $t \in [0, T]$ it holds that

$$\|\nabla_x u(x, t) - \nabla_x V(x, t)\| \leq \varepsilon(1 + \|x\|^\gamma) \quad (23)$$

and that

$$\|\nabla_x \phi(x, t)\| \leq \kappa(1 + \|x\|^\gamma). \quad (24)$$

Then the variance can be bounded by

$$\mathbb{V} \left[\left. \frac{\delta}{\delta u} \widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}(u; \phi) \right] \leq \frac{C\varepsilon^2 \kappa^2}{K}, \quad (25)$$

where $C \in \mathbb{R}_{>0}$ is a constant that depends only on γ, ξ, τ , and the given PDE in (1).

Proof. See Appendix B. \square

Remark 2.8. We remark that the simultaneous approximation of a function and its gradient by a neural network, as required in (22) and (23), has, for instance, been considered in Berner et al. (2019). In Proposition 2.7 we can see that it is actually sufficient to approximate the gradient $\nabla_x V$ (which only implies the approximation of V up to a constant) in order to reach zero variance. We also note that, for u far away from the solution V , the term S_u could potentially also increase the estimator variance, see, e.g., Figure 14 in the appendix.

2.2. Algorithmic Details and Neural Network Approximations

So far our analysis has been agnostic of the choice of the approximating function space \mathcal{U} . In practice, we usually rely on functions represented by a neural network $\Phi_\theta: \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ with parameters $\theta \in \mathbb{R}^p$. The training of such a network is then conducted by minimizing a suitable estimated loss $\widehat{\mathcal{L}}^{(K)}$ using some variant of gradient descent. In its simplest form, this corresponds to the update rule

$$\theta^{(m+1)} = \theta^{(m)} - \lambda \nabla_\theta \widehat{\mathcal{L}}^{(K)}(\Phi_{\theta^{(m)}}), \quad (26)$$

where $\lambda \in \mathbb{R}_{>0}$ is a learning rate and $\theta^{(0)} \in \mathbb{R}^p$ is a random initialization. Let us refer to Algorithm 1 in Appendix E for further details.

Relating to the Gateaux derivatives from the previous section, let us note that we are particularly interested in the directions⁵ $\phi = \partial_{\theta_i} \Phi_\theta$ for $i \in \{1, \dots, p\}$. This choice is motivated by (26) and the chain rule of the Gateaux derivative, which, under suitable assumptions, states that

$$\partial_{\theta_i} \widehat{\mathcal{L}}^{(K)}(\Phi_\theta) = \left. \frac{\delta}{\delta u} \right|_{u=\Phi_\theta} \widehat{\mathcal{L}}^{(K)}(u; \partial_{\theta_i} \Phi_\theta). \quad (27)$$

Now, in order to elaborate on some computational aspects, let us state the actual gradients of the corresponding losses. For a neural network $u = \Phi_\theta$, we compute

$$\nabla_\theta \widehat{\mathcal{L}}_{\text{FK}}^{(K)}(u) = -\frac{2}{K} \sum_{k=1}^K \widehat{\Delta}_u^{(k)} \nabla_\theta u, \quad (28a)$$

$$\nabla_\theta \widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}(u) = -\frac{2}{K} \sum_{k=1}^K e^{(k)} \left(\nabla_\theta u + \nabla_\theta \widehat{S}_u^{(k)} \right), \quad (28b)$$

where $e^{(k)} := \widehat{\Delta}_u^{(k)} - \widehat{S}_u^{(k)}$ is the sample-wise error.

From a computational point of view, note that the derivative of $\widehat{\mathcal{L}}_{\text{FK}}^{(K)}$ as displayed in (28a) only relies on a single gradient computation of the approximating function u (w.r.t. the model parameters θ), whereas the derivative of $\widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}$ as displayed in (28b) needs evaluations of $\nabla_\theta s_j^{(k)}$, see (16) and (17). This might be costly for automatic differentiation (autodiff) tools and we will suggest computational speed-ups and memory-efficient versions in the next section.

Further, note that due to the term $\widehat{S}_u^{(k)}$ the loss $\widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}$ always needs to rely on discretized SDEs, whereas $\widehat{\mathcal{L}}_{\text{FK}}^{(K)}$ can be computed without SDE simulations whenever explicit solutions of the stochastic processes are available (which, admittedly, is often not the case in practice – see, however, the examples in Sections 3.1 and 3.2).

⁵We assume that the functions $\Phi_\theta, \theta \in \mathbb{R}^p$, as well as all partial derivatives $\partial_{\theta_i} \Phi_\theta$ lie in \mathcal{U} .

2.3. Further Computational Adjustments

For the loss $\mathcal{L}_{\text{BSDE}}$ as defined in (11) we can make further adjustments that will turn out to be relevant from a computational point of view. In particular, let us suggest two alternative versions that are based on the fact that the function u in (11) appears in two instances, which can be treated differently. One idea is to differentiate only with respect to one of its appearances. To this end, let us define the loss

$$\mathcal{L}_{\text{BSDE}}^{\text{detach}}(u, w) := \mathbb{E} \left[(\Delta_u - S_w)^2 \right], \quad (29)$$

now depending on two arguments. We then compute derivatives according to

$$\left. \frac{\delta}{\delta u} \right|_{w=u} \mathcal{L}_{\text{BSDE}}^{\text{detach}}(u, w), \quad (30)$$

where in automatic differentiation tools setting $w = u$ after computing the derivative is usually obtained by detaching w from the computational graph⁶. One can readily check that $\mathcal{L}_{\text{BSDE}}^{\text{detach}}$ is a valid loss for problem (1) that keeps the robustness property stated in Propositions 2.4 and 2.6 intact. At the same time we expect a great reduction of the computational effort since the second-order derivatives of the form $\nabla_\theta \partial_{x_i} u$, appearing in $\nabla_\theta \widehat{S}_u^{(k)}$, do not need to be computed anymore – in particular note that multiple such derivatives would need to be evaluated in actual implementations due to the discretization of the integral, as stated in (16b). We refer to Section 3 where we can indeed demonstrate substantial computational improvements in several settings.

Another version of $\mathcal{L}_{\text{BSDE}}$ relies on explicitly modelling the gradient of the approximating function u by an extra function, see also Zhou et al. (2021). We can therefore define the loss

$$\mathcal{L}_{\text{BSDE}}^{\text{grad}}(u, r) := \mathbb{E} \left[(\Delta_u - \widetilde{S}_r)^2 \right], \quad (31)$$

where now

$$\widetilde{S}_r = \int_\tau^T (\sigma^\top r)(X_s, s) \cdot dW_s. \quad (32)$$

The additional function $r: \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}^d$ can be modelled with a separate neural network. The loss (31) is then minimized with respect to the parameters of the functions u and r simultaneously, which again avoids the computation of second-order derivatives.

Finally, let us introduce another viable option for decreasing computational resources, which is essentially a more memory-efficient way of computing the derivative formula stated in (28b). Usually, one performs a forward pass to

⁶In PyTorch and TensorFlow this can be achieved by the `detach` and `stop_gradient` operations.

compute $\widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}(u)$ and then uses automatic differentiation to obtain the gradient $\nabla_{\theta} \widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}(u)$. However, this requires to keep track of all the discretization steps for the stochastic integral $(s_j^{(k)})_{j=1}^{J^{(k)}}$, as they all depend on θ . This leads to a GPU memory consumption, which scales linearly in the number of steps $J^{(k)}$ and, already for small K , exceeds common capacities, see Figures 5 and 7 in the appendix.

To circumvent this issue, we observe that the derivative in (28b) can be decomposed into

$$G_1 := -\frac{2}{K} \sum_{k=1}^K e^{(k)} \nabla_{\theta} u, \quad (33)$$

which equals the derivative of the loss $\mathcal{L}_{\text{BSDE}}^{\text{detach}}$, and

$$\sum_{j=1}^J G_{2,j} := \sum_{j=1}^J -\frac{2}{K} \sum_{k=1}^K e^{(k)} \nabla_{\theta} s_j^{(k)}, \quad (34)$$

where $J = \max_{k=1}^K J^{(k)}$ and $s_j^{(k)} = 0$ for $j > J^{(k)}$. Therefore, we can first compute and cache the errors $(e^{(k)})_{k=1}^K$ and the gradients G_1 . Then, we repeat the same simulation, i.e., use the same realization of $(\xi^{(k)}, \tau^{(k)}, W^{(k)})$, to compute and accumulate the gradients $(G_{2,j})_{j=1}^J$ on-the-fly. In doing so, automatic differentiation only needs to track a single discretization step $s_j^{(k)}$. This keeps the memory footprint independent of the number of steps $J^{(k)}$ or, equivalently, the step-size $t_{j+1}^{(k)} - t_j^{(k)}$ – however, at the cost of an increased computational time⁷. Of course, the same trick can also be applied to the loss $\mathcal{L}_{\text{BSDE}}^{\text{grad}}$. Even though from a mathematical perspective the losses are still the same as $\mathcal{L}_{\text{BSDE}}$ and $\mathcal{L}_{\text{BSDE}}^{\text{grad}}$, we give them the new names

$$\mathcal{L}_{\text{BSDE}, \text{eff}} \quad \text{and} \quad \mathcal{L}_{\text{BSDE}, \text{eff}}^{\text{grad}} \quad (35)$$

in order to relate to their significant computational improvements, which we will illustrate in the next section.

3. Numerical Examples

In this section we aim to illustrate our theoretical results on three high-dimensional PDEs⁸. Specifically, we consider a heat equation with paraboloid terminal condition, a HJB equation arising in molecular dynamics, and a Black-Scholes model with correlated noise. The respective reference solutions are given as a closed-form expression, as a

⁷For most practical use-cases one is interested in obtaining the most accurate solution and there is only little restriction on the time for training, which needs to be done once per PDE. However, for time-critical applications, we present the performance w.r.t. the training time in Figures 11 and 12 in the appendix.

⁸The associated code can be found at https://github.com/juliusberner/robust_kolmogorov.

tensor-product of one-dimensional approximations obtained by finite-difference methods, and as a Monte-Carlo approximation using the Feynman-Kac formula, see Appendix C.

In all of our numerical experiments we use comparable setups which are detailed in Appendix E. We emphasize that, instead of using a train/val/test split on a given finite data set, our setting allows us to simulate new i.i.d. samples $(\xi^{(k)}, \tau^{(k)}, W^{(k)})$ for each batch on demand during training. To evaluate our methods, we compute the MSE between the neural network approximation and the reference solution at points (x, t) which are sampled independently of the training data, i.e., on previously unseen points, see Appendix E.

We will show that all considered losses are viable in the sense that they yield appropriate approximations of solutions to the associated PDEs. At the same time we will see substantial performance differences, which can be attributed to the different variance properties outlined before.

3.1. Heat Equation

Let us first consider a version of the heat equation with paraboloid terminal condition by setting $b(x, t) = 0$, $\sigma(x, t) = \bar{\sigma} \in \mathbb{R}^{d \times d}$, and $g(x) = \|x\|^2$. This allows for explicit solutions given by

$$V(x, t) = \|x\|^2 + \text{Trace}(\bar{\sigma} \bar{\sigma}^{\top})(T - t) \quad (36)$$

and

$$X_s = \xi + \bar{\sigma}(W_s - W_{\tau}), \quad (37)$$

see also Beck et al. (2021). We analyze the special case $\bar{\sigma} = \nu \text{Id}$ with $\nu \in \mathbb{R}$, where we have that

$$\int_{\tau}^T \bar{\sigma}^{\top} \nabla_x V(X_s, s) \cdot dW_s = 2\nu \int_{\tau}^T X_s \cdot dW_s. \quad (38)$$

Using the fact that $\int_0^T B_s \cdot dB_s = \frac{1}{2}(\|B_T\|^2 - d\mathcal{T})$, where $\mathcal{T} = T - \tau$ and $B_s = W_{\tau+s} - W_{\tau}$, the stochastic integral in (38) evaluates to the following closed-form formula

$$2\nu \xi \cdot (W_T - W_{\tau}) + \nu^2 (\|W_T - W_{\tau}\|^2 - d(T - \tau)). \quad (39)$$

This expression shows how the quantities in Propositions 2.4 and 2.6 may behave. For a fair comparison, however, we do not make use of this explicit representation.

For our experiments, we consider $d = 50$, $T = 1$, $\nu = 0.5$, $\xi \sim \text{Unif}([-0.5, 0.5]^d)$, and $\tau \sim \text{Unif}([0, 1])$. For a fixed batch size all of our proposed methods significantly outperform the standard approach of minimizing \mathcal{L}_{FK} , see Figure 1. For larger batch sizes the losses $\mathcal{L}_{\text{BSDE}}$ and $\mathcal{L}_{\text{BSDE}}^{\text{grad}}$ run out of our memory limit of 8 GiB and the performance of \mathcal{L}_{FK} improves. However, when using the efficient versions or $\mathcal{L}_{\text{BSDE}}^{\text{detach}}$, we still outperform the baseline by a substantial margin. As suggested by our theory, the performance gains




Figure 4. Standard deviation of the gradient estimators (maximum over directions $\phi = \partial_{\theta_i} \Phi_\theta$) and MSE of the gradients (on evaluation data) in the setup of Figure 1, see also Appendix E. In line with Proposition 2.7 and Remark 2.8, the standard deviations of the loss and gradient estimators are smaller for the BSDE-based losses, especially when the gradient $\nabla_x V$ is well approximated.

are based on the variance reducing effect of the stochastic integral, see Figures 1 and 4. While the gradient $\nabla_x V$ can be best approximated by $\mathcal{L}_{\text{BSDE,eff}}^{\text{grad}}$ (see also Figure 8 in the appendix), which reduces the variance according to Proposition 2.7, the loss $\mathcal{L}_{\text{BSDE,eff}}$ still performs better in approximating the solution V . This can be motivated by the fact that in case of $\mathcal{L}_{\text{BSDE,eff}}^{\text{grad}}$ one network learns the gradient and another one independently learns the solution. In contrast, in the case of $\mathcal{L}_{\text{BSDE,eff}}$, a single network is trained to simultaneously approximate V and its gradient $\nabla_x V$. Note that these observations are consistent across various diffusivities ν , see Figures 9 and 10.

3.2. Black–Scholes Model with Correlated Noise

As a second example we consider a version of the celebrated Black–Scholes model from financial engineering (Black & Scholes, 1973), given by

$$\sigma(x, t) = \text{diag}(\beta_1 x_1, \dots, \beta_d x_d) \bar{\sigma}, \quad b(x, t) = \bar{b}x, \quad (40)$$

with $\beta \in \mathbb{R}^d$, $\bar{\sigma} \in \mathbb{R}^{d \times d}$, $\bar{b} \in \mathbb{R}$, and

$$g(x) = \max \left\{ 0, \kappa - \min_{i=1}^d x_i \right\}. \quad (41)$$

It models a rainbow European put option, giving its holder the right to sell the minimum of underlying assets at the strike price $\kappa \in \mathbb{R}_{>0}$ at time T .

For the experiments we choose $d = 50$, $T = 1$, $\bar{b} = -0.05$, $\kappa = 5.5$, $\beta = (0.1 + i/(2d))_{i=1}^d$,

$$\xi \sim \text{Unif}([4.5, 5.5]^d), \quad \tau \sim \text{Unif}([0, 1]), \quad (42)$$

and $\bar{\sigma}$ to be the lower triangular matrix arising from the Cholesky decomposition $\bar{\sigma} \bar{\sigma}^\top = Q$ with

$$Q_{i,j} = 0.5(1 + \delta_{i,j}), \quad i, j \in \{1, \dots, d\}, \quad (43)$$

where $\delta_{i,j}$ denotes the Kronecker delta. There does not exist a closed form solution V , however, one can use Monte Carlo




Figure 5. Estimator standard deviation, gradient estimator standard deviation (maximum over directions $\phi = \partial_{\theta_i} \Phi_\theta$), and MSE (on evaluation data) after $30k$ gradient steps for the Black–Scholes model in Section 3.2. The last plot shows that $\mathcal{L}_{\text{BSDE}}$ and $\mathcal{L}_{\text{BSDE}}^{\text{grad}}$ cannot be used for batch size $K = 1024$ and larger as it would exceed the GPU memory limit of 8 GiB. The same holds for the other losses at batch size $K = 524288$.

sampling and the representation

$$(X_s)_i = \xi_i e^{\left(\bar{b} - \frac{\|\beta_i \sigma_i\|^2}{2}\right)(s-\tau) + \beta_i \sigma_i \cdot (W_s - W_\tau)}, \quad (44)$$

where σ_i is the i -th row of $\bar{\sigma}$, in order to evaluate the solution pointwise, see Beck et al. (2021). Note that, in case of the loss \mathcal{L}_{FK} , such a closed-form expression could be used instead of the Euler–Maruyama scheme in order to speed up training. However, for most settings such a representation is not available and, for the sake of a fair comparison, we did not make use of it during training.

Note that for this example the coefficient functions σ and b depend on x , which makes the discretization of the SDE more delicate. The loss $\mathcal{L}_{\text{BSDE}}^{\text{detach}}$ can be sensitive to initial performance and performs suboptimally in this example, see Remark 2.8 and Figure 14. The other losses again underline our theoretical results and, in particular, substantially outperform the loss \mathcal{L}_{FK} , see Figure 5. As suggested by our results, a larger batch size generally improves the robustness. However, the efficient versions of $\mathcal{L}_{\text{BSDE}}$ and $\mathcal{L}_{\text{BSDE}}^{\text{grad}}$ do not need significantly more GPU memory than \mathcal{L}_{FK} such that they can be used with the same maximal batch size.

3.3. Hamilton–Jacobi–Bellman Equation

Finally, let us consider the nonlinear Hamilton–Jacobi–Bellman equation

$$(\partial_t + L)\tilde{V}(x, t) + f(x, t) = \frac{1}{2} \|(\sigma^\top \nabla_x \tilde{V})(x, t)\|^2 \quad (45)$$

for $(x, t) \in \mathbb{R}^d \times [0, T]$, with terminal condition

$$\tilde{V}(x, T) = \tilde{g}(x) \quad (46)$$

for $x \in \mathbb{R}^d$. This prominent PDE from control theory corresponds to the controlled stochastic process

$$dX_s^v = (b(X_s^v, s) + (\sigma v)(X_s^v, s)) ds + \sigma(X_s^v, s) dW_s,$$

where the control $v \in C(\mathbb{R}^d \times [0, T], \mathbb{R}^d)$ can be thought of as a steering force to be chosen so as to minimize a given cost functional

$$J(v) = \mathbb{E} \left[\int_0^T \left(f_s + \frac{1}{2} \|v_s\|^2 \right) ds + \tilde{g}(X_T^v) \right], \quad (47)$$

with the short-hands $f_s := f(X_s^v, s)$ and $v_s := v(X_s^v, s)$. In the above, $f \in C(\mathbb{R}^d \times [0, T], \mathbb{R})$ represents running costs (in addition to the squared costs on the control v) and $\tilde{g} \in C(\mathbb{R}^d, \mathbb{R})$ specifies the terminal costs. It turns out that one can recover the optimal control v^* that minimizes the costs (47) from the solution to the PDE in (45) by the relation $v^* = -\sigma^\top \nabla_x \tilde{V}$.

Now, with the transformation⁹ $V = \exp(-\tilde{V})$, we can convert the nonlinear PDE in (45) to a linear one of the form (1) (or, to be precise, of the form (69)) with boundary condition $g = \exp(-\tilde{g})$, which allows us to apply Algorithm 1, see Lemma G.1 for the details.

For our experiments, we consider a problem that is prominent in molecular dynamics and has been suggested in Nüsken & Richter (2021b). We define the drift to be $b = -\nabla_x \Psi$, with $\Psi(x) = \kappa \sum_{i=1}^d (x_i^2 - 1)^2$ being a potential of double well type. In applications, the potential represents the energy associated to a system of atoms (i.e. molecule) and one is often interested in how molecular configurations change over time. This is then related to transition paths between metastable states of the stochastic dynamics, whose sampling poses formidable computational challenges, see e.g. Hartmann & Richter (2021).

For an example, let us define a terminal value that is supported in one of the 2^d minima of Ψ , namely $g(x) = \exp(-\eta \sum_{i=1}^d (x_i - 1)^2)$ and set $f = 0$. We choose $\eta = 0.04$ and $\kappa = 0.1$ in dimension $d = 10$ as well as $\sigma = \text{Id}$. The left panel of Figure 6 shows the original potential Ψ as well as a tilted version that can be derived from the solution of the HJB equation (45) via $\Psi^* = \Psi + \sigma \sigma^\top \tilde{V}$, computed with a finite difference reference method as well as with our algorithm. We can see that both solutions agree for all of our considered methods. However, note that for the optimal control functions the BSDE-based losses are again superior to \mathcal{L}_{FK} , as displayed in the right panel. This is in line with the more accurate approximation of the gradi-

⁹This transformation is also known as *Hopf-Cole transformation*, see, e.g., Section 4.4.1 in Evans (2010). See also Hartmann et al. (2017) for a discussion on related applications in importance sampling of stochastic processes relevant in computational statistical physics.




Figure 6. For a HJB equation with double well potential we display the tilted version of the potential Ψ^* evaluated at $x = (\alpha, \dots, \alpha)^\top$ and $t = 0.75$, which is approximated well by all our methods (after 30k gradient steps with batch size 1024). However, the approximation of the first component of the control $\sigma^\top \nabla_x \log V$ depicted in the right panel is better for BSDE-based losses.

ent, which is demonstrated by a performance comparison in Figure 15 in the appendix.

4. Conclusion

The most prominent method for solving general high-dimensional linear PDEs of Kolmogorov type is based on minimizing a variational formulation inspired by the Feynman-Kac formula over a class of neural networks. Building upon results on BSDEs and control variates for Monte Carlo estimators, we suggest an alternative formulation and prove that the corresponding loss estimator enjoys lower variance. Lending to the notion of Gateaux derivatives, we can further show that this leads to gradient estimators with lower variance, which is crucial for gradient-based optimization. Importantly, we also develop a novel, more efficient estimator that is adapted to current deep learning frameworks and show that resulting methods yield significant performance gains. It is up to future research to exploit such efficient estimators for boundary value problems, see Appendix C, and non-linear PDEs, which can naturally be approached with BSDE-based losses, see Appendix D.

In summary, we suspect that the techniques used in this paper can be transferred to a number of related deep learning methods that rely on Monte Carlo sampling. Providing important theoretical guarantees and performance improvements, our work is a fundamental step in developing robust and reliable algorithms for the solution of high-dimensional PDEs.

Acknowledgements

The research of Lorenz Richter has been funded by Deutsche Forschungsgemeinschaft (DFG) through the grant CRC 1114 “Scaling Cascades in Complex Systems” (project A05, project number 235221301). The research of Julius Berner was supported by the Austrian Science Fund (FWF) under grant I3403-N32. The computational results have been achieved in part using the Vienna Scientific Cluster (VSC).

References

- Ash, R. B. and Doléans-Dade, C. *Probability and measure theory*. Academic press, 2000.
- Baldi, P. *Stochastic Calculus: An Introduction Through Theory and Exercises*. Universitext. Springer International Publishing, 2017.
- Beck, C., E. W., and Jentzen, A. Machine learning approximation algorithms for high-dimensional fully nonlinear partial differential equations and second-order backward stochastic differential equations. *Journal of Nonlinear Science*, 29(4):1563–1619, 2019.
- Beck, C., Becker, S., Grohs, P., Jaafari, N., and Jentzen, A. Solving the Kolmogorov PDE by means of deep learning. *Journal of Scientific Computing*, 88(3):1–28, 2021.
- Berner, J., Elbrächter, D., Grohs, P., and Jentzen, A. Towards a regularity theory for ReLU networks—chain rule and global error estimates. In *13th International conference on Sampling Theory and Applications*, pp. 1–5. IEEE, 2019.
- Berner, J., Dablander, M., and Grohs, P. Numerically solving parametric families of high-dimensional Kolmogorov partial differential equations via deep learning. In *Advances in Neural Information Processing Systems*, pp. 16615–16627, 2020a.
- Berner, J., Grohs, P., and Jentzen, A. Analysis of the generalization error: Empirical risk minimization over deep artificial neural networks overcomes the curse of dimensionality in the numerical approximation of Black–Scholes partial differential equations. *SIAM Journal on Mathematics of Data Science*, 2(3):631–657, 2020b.
- Berner, J., Grohs, P., Kutyniok, G., and Petersen, P. The modern mathematics of deep learning. *arXiv preprint arXiv:2105.04026*, 2021.
- Bismut, J.-M. Conjugate convex functions in optimal stochastic control. *Journal of Mathematical Analysis and Applications*, 44(2):384–404, 1973.
- Black, F. and Scholes, M. The pricing of options and corporate liabilities. *Journal of political economy*, 81(3):637–654, 1973.
- Da Prato, G. and Zabczyk, J. *Stochastic equations in infinite dimensions*. Cambridge university press, 2014.
- E, W., Han, J., and Jentzen, A. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Communications in Mathematics and Statistics*, 5(4):349–380, 2017.
- Ekström, E. and Tysk, J. The black–scholes equation in stochastic volatility models. *Journal of Mathematical Analysis and Applications*, 368(2):498 – 507, 2010.
- Evans, L. C. *Partial differential equations*. American Mathematical Society, Providence, R.I., 2010.
- Gall, J. *Brownian Motion, Martingales, and Stochastic Calculus*. Graduate Texts in Mathematics. Springer International Publishing, 2016.
- Gobet, E. *Monte-Carlo methods and stochastic processes: from linear to non-linear*. CRC Press, 2016.
- Grohs, P. and Herrmann, L. Deep neural network approximation for high-dimensional elliptic PDEs with boundary conditions. *arXiv preprint arXiv:2007.05384*, 2020.
- Hairer, M., Hutzenthaler, M., and Jentzen, A. Loss of regularity for Kolmogorov equations. *The Annals of Probability*, 43(2):468–527, 2015.
- Han, J., Jentzen, A., and Weinan, E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- Hartmann, C. and Richter, L. Nonasymptotic bounds for suboptimal importance sampling. *arXiv preprint arXiv:2102.09606*, 2021.
- Hartmann, C., Richter, L., Schütte, C., and Zhang, W. Variational characterization of free energy: Theory and algorithms. *Entropy*, 19(11):626, 2017.
- Hartmann, C., Kebiri, O., Neureither, L., and Richter, L. Variational approach to rare event simulation using least-squares regression. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(6):063107, 2019.
- Hendrycks, D. and Gimpel, K. Gaussian error linear units (GELUs). *arXiv preprint arXiv:1606.08415*, 2016.
- Huang, C.-W., Lim, J. H., and Courville, A. C. A variational perspective on diffusion-based generative models and score matching. *Advances in Neural Information Processing Systems*, 34, 2021.
- Jentzen, A., Salimova, D., and Welti, T. A proof that deep artificial neural networks overcome the curse of dimensionality in the numerical approximation of Kolmogorov partial differential equations with constant diffusion and nonlinear drift coefficients. *arXiv preprint arXiv:1809.07321*, 2018.
- Karatzas, I. and Shreve, S. E. *Brownian Motion and Stochastic Calculus*. Springer, 1998.

- Kiumarsi, B., Vamvoudakis, K. G., Modares, H., and Lewis, F. L. Optimal and autonomous control using reinforcement learning: A survey. *IEEE transactions on neural networks and learning systems*, 29(6):2042–2062, 2017.
- Klenke, A. *Probability Theory: A Comprehensive Course*. Universitext. Springer London, 2013.
- Kloeden, P. E. and Platen, E. Stochastic differential equations. In *Numerical Solution of Stochastic Differential Equations*, pp. 103–160. Springer, 1992.
- Lelievre, T. and Stoltz, G. Partial differential equations and stochastic methods in molecular dynamics. *Acta Numerica*, 25:681, 2016.
- Nüsken, N. and Richter, L. Interpolating between BSDEs and PINNs: deep learning for elliptic and parabolic boundary value problems. *arXiv preprint arXiv:2112.03749*, 2021a.
- Nüsken, N. and Richter, L. Solving high-dimensional Hamilton–Jacobi–Bellman PDEs using neural networks: perspectives from the theory of controlled diffusions and measures on path space. *Partial Differential Equations and Applications*, 2(4):1–48, 2021b.
- Øksendal, B. Stochastic differential equations. In *Stochastic differential equations*, pp. 65–84. Springer, 2003.
- Pardoux, É. Backward stochastic differential equations and viscosity solutions of systems of semilinear parabolic and elliptic PDEs of second order. In *Stochastic Analysis and Related Topics VI*, pp. 79–127. Springer, 1998.
- Pardoux, E. and Peng, S. Adapted solution of a backward stochastic differential equation. *Systems & Control Letters*, 14(1):55–61, 1990.
- Pascucci, A. Kolmogorov equations in physics and in finance. In *Elliptic and parabolic problems*, pp. 353–364. Springer, 2005.
- Pavliotis, G. A. *Stochastic processes and applications: diffusion processes, the Fokker-Planck and Langevin equations*, volume 60. Springer, 2014.
- Pfau, D., Spencer, J. S., Matthews, A. G., and Foulkes, W. M. C. Ab initio solution of the many-electron Schrödinger equation with deep neural networks. *Physical Review Research*, 2(3):033429, 2020.
- Pham, H. *Continuous-time stochastic control and optimization with financial applications*, volume 61. Springer Science & Business Media, 2009.
- Pironneau, O. and Achdou, Y. Partial differential equations for option pricing. *Handbook of Numerical Analysis*, 15: 369–495, 2009.
- Raissi, M., Perdikaris, P., and Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- Reisinger, C. and Zhang, Y. Rectified deep neural networks overcome the curse of dimensionality for nonsmooth value functions in zero-sum games of nonlinear stiff systems. *arXiv:1903.06652*, 2019.
- Richter, L. *Solving high-dimensional PDEs, approximation of path space measures and importance sampling of diffusions*. PhD thesis, BTU Cottbus-Senftenberg, 2021.
- Robert, C. P. and Casella, G. *Monte Carlo statistical methods*, volume 2. Springer, 2004.
- Sirignano, J. and Spiliopoulos, K. DGM: A deep learning algorithm for solving partial differential equations. *Journal of Computational Physics*, 375:1339–1364, 2018.
- Theodorou, E., Buchli, J., and Schaal, S. A generalized path integral control approach to reinforcement learning. *The Journal of Machine Learning Research*, 11:3137–3181, 2010.
- Vidales, M. S., Siska, D., and Szpruch, L. Unbiased deep solvers for parametric PDEs. *arXiv preprint arXiv:1810.05094*, 2018.
- Widder, D. V. *The heat equation*, volume 67. Academic Press, 1976.
- Zhang, J. *Backward stochastic differential equations*. Springer, 2017.
- Zhou, M., Han, J., and Lu, J. Actor-critic method for high dimensional static Hamilton–Jacobi–Bellman partial differential equations based on neural networks. *arXiv preprint arXiv:2102.11379*, 2021.

A. Notation

We write $C^{2,1}(\mathbb{R}^d \times [0, T], \mathbb{R})$ for the space of functions $(x, t) \mapsto f(x, t)$ which are twice continuously differentiable in the spatial coordinate $x \in \mathbb{R}^d$ and once continuously differentiable in the time coordinate $t \in [0, T]$. We further denote the partial derivatives (w.r.t. to the i -th spatial coordinate x_i and time coordinate t) and the gradient of f (w.r.t. to the spatial coordinate x) by $\partial_{x_i} f$, $\partial_t f$, and $\nabla_x f = (\partial_{x_1} f, \dots, \partial_{x_d} f)^\top$, respectively. We say that functions $f \in C(\mathbb{R}^d \times [0, T], \mathbb{R})$ and $g \in C(\mathbb{R}^d, \mathbb{R})$ are *at most polynomially growing* if there exist constants $c, \lambda \in \mathbb{R}_{>0}$ such that for every $(x, t) \in \mathbb{R}^d \times [0, T]$ it holds that

$$|g(x)| \leq c(1 + \|x\|^\lambda), \quad |f(x, t)| \leq c(1 + \|x\|^\lambda). \quad (48)$$

Finally, we write $\|\cdot\|$ for the Euclidean norm and $\|\cdot\|_F$ for the Frobenius norm.

B. Assumptions and Proofs

We usually operate under the following assumption which guarantees a unique strong solution to the SDE in (6) as well as the existence of a unique strong solution $V \in C^{2,1}(\mathbb{R}^d \times [0, T], \mathbb{R})$ to the PDE in (1) which can be expressed as $V(x, t) = \mathbb{E}[g(X_T) | X_t = x]$ as, for instance, stated in (13), see Baldi (2017, Theorem 10.6), Karatzas & Shreve (1998, Remark 5.7.8), Øksendal (2003, Theorem 5.2.1), and Pavliotis (2014, Theorem 2.1).

Assumption 1 (Conditions on b , σ , and g). *We assume that $g \in C^2(\mathbb{R}^d, \mathbb{R})$ is an at most polynomially growing function. Further, we assume that the coefficient functions $\sigma \in C(\mathbb{R}^d \times [0, T], \mathbb{R}^{d \times d})$ and $b \in C(\mathbb{R}^d \times [0, T], \mathbb{R}^d)$ satisfy the following properties: There exist constants $c_1, c_2, c_3, c_4 \in \mathbb{R}_{>0}$ such that*

$$\begin{aligned} \|b(x, t)\| + \|\sigma \sigma^\top(x, t)\|_F &\leq c_1, \quad (\text{boundedness}) \\ \|\sigma(x, t) - \sigma(y, t)\|_F &\leq c_2 \|x - y\|, \quad (\text{Lipschitz continuity}) \\ \|b(x, t) - b(y, t)\| &\leq c_3 \|x - y\|, \quad (\text{Lipschitz continuity}) \\ \eta \cdot (\sigma \sigma^\top)(x, t) \eta &\geq c_4 \|\eta\|^2, \quad (\text{uniform ellipticity}) \end{aligned}$$

for all $x, y, \eta \in \mathbb{R}^d$ and $t \in [0, T]$.

Proof of Proposition 2.2. We will use the fact that $\Delta_V = S_V$ almost surely, see Lemma 2.1. This allows us to rewrite

$$\begin{aligned} \mathcal{L}_{\text{BSDE}}(u) &= \mathbb{E} \left[(\Delta_u + V(\xi, \tau) - V(\xi, \tau) - S_u)^2 \right] \\ &= \mathbb{E} \left[(V(\xi, \tau) - u(\xi, \tau) + S_V - S_u)^2 \right] \\ &= \mathbb{E} \left[(V(\xi, \tau) - u(\xi, \tau))^2 \right] + \mathbb{E} [S_V^2 - u]. \end{aligned}$$

In the last step we used the tower property of the conditional expectation and the vanishing expectation of the stochastic

integrals S_V and S_u to show that the cross-term has zero expectation, namely

$$\begin{aligned} \mathbb{E} [(V(\xi, \tau) - u(\xi, \tau)) (S_V - S_u)] \\ = \mathbb{E} [(V(\xi, \tau) - u(\xi, \tau)) \mathbb{E} [S_V - S_u | (\xi, \tau)]] = 0. \end{aligned}$$

For the loss \mathcal{L}_{FK} we can similarly write

$$\mathcal{L}_{\text{FK}}(u) = \mathbb{E} \left[(V(\xi, \tau) - u(\xi, \tau) + \Delta_V)^2 \right] \quad (49a)$$

$$= \mathbb{E} \left[(V(\xi, \tau) - u(\xi, \tau))^2 \right] + \mathbb{E} [\Delta_V^2], \quad (49b)$$

where the cross-term has again zero expectation. More precisely, the tower property of the conditional expectation and Lemma 2.1 imply that

$$\mathbb{E} [(V(\xi, \tau) - u(\xi, \tau)) \Delta_V] \quad (50a)$$

$$= \mathbb{E} [(V(\xi, \tau) - u(\xi, \tau)) \mathbb{E} [\Delta_V | (\xi, \tau)]] \quad (50b)$$

$$= \mathbb{E} [(V(\xi, \tau) - u(\xi, \tau)) \mathbb{E} [S_V | (\xi, \tau)]] = 0. \quad (50c)$$

This proves Proposition 2.2. \square

Proof of Proposition 2.4. Let us recall Lemma 2.1 which states that $\Delta_V = S_V$ almost surely. With the mutual independence of $(\widehat{\Delta}_u^{(k)})_{k=1}^K$ and $(\widehat{S}_u^{(k)})_{k=1}^K$ we now readily obtain that

$$\mathbb{E} \left[\widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}(V) \right] = \mathbb{E} \left[(\Delta_V - S_V)^2 \right] = 0, \quad (51)$$

and

$$\mathbb{V} \left[\widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}(V) \right] = \frac{1}{K} \mathbb{V} \left[(\Delta_V - S_V)^2 \right] = 0. \quad (52)$$

For the loss \mathcal{L}_{FK} , it holds that

$$\mathbb{E} \left[\widehat{\mathcal{L}}_{\text{FK}}^{(K)}(V) \right] = \mathbb{E} [\Delta_V^2] = \mathbb{V} [S_V] \quad (53)$$

and

$$\mathbb{V} \left[\widehat{\mathcal{L}}_{\text{FK}}^{(K)}(V) \right] = \frac{1}{K} \mathbb{V} [\Delta_V^2] = \frac{1}{K} \mathbb{V} [S_V^2], \quad (54)$$

which proves the claim. \square

Proof of Proposition 2.6. (i) It holds that

$$\frac{\delta}{\delta u} \widehat{\mathcal{L}}_{\text{FK}}^{(K)}(u; \phi) = \frac{d}{d\varepsilon} \Big|_{\varepsilon=0} \widehat{\mathcal{L}}_{\text{FK}}^{(K)}(u + \varepsilon \phi) \quad (55a)$$

$$= \frac{d}{d\varepsilon} \Big|_{\varepsilon=0} \frac{1}{K} \sum_{k=1}^K \left(\Delta_{u+\varepsilon \phi}^{(k)} \right)^2 \quad (55b)$$

$$= -\frac{2}{K} \sum_{k=1}^K \Delta_u^{(k)} \phi(\xi^{(k)}, \tau^{(k)}). \quad (55c)$$

Using the mutual independence of our samples this implies that

$$\mathbb{W} \left[\frac{\delta}{\delta u} \widehat{\mathcal{L}}_{\text{FK}}^{(K)}(u; \phi) \right] = \frac{4}{K} \mathbb{W} [\Delta_u \phi(\xi, \tau)]. \quad (56)$$

Now, setting $u = V$, Lemma 2.1 shows that the variance of the Gateaux derivative at V satisfies

$$\mathbb{W} \left[\frac{\delta}{\delta u} \Big|_{u=V} \widehat{\mathcal{L}}_{\text{FK}}^{(K)}(u; \phi) \right] = \frac{4}{K} \mathbb{W} [S_V \phi(\xi, \tau)]. \quad (57)$$

By the Itô isometry, this evaluates to

$$\frac{4}{K} \mathbb{E} \left[\left(\int_{\tau}^T \|\sigma^\top \nabla_x V(X_s, s)\|^2 ds \right) \phi(\xi, \tau) \right], \quad (58)$$

which, in general, is non-zero for an arbitrary $\phi \in \mathcal{U}$.

(ii) By definition, we have that

$$\widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}(u + \varepsilon \phi) = \frac{1}{K} \sum_{k=1}^K \left(\Delta_{u+\varepsilon \phi}^{(k)} - S_{u+\varepsilon \phi}^{(k)} \right)^2. \quad (59)$$

The Gateaux derivative

$$\frac{\delta}{\delta u} \widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}(u; \phi) = \frac{d}{d\varepsilon} \Big|_{\varepsilon=0} \widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}(u + \varepsilon \phi) \quad (60)$$

thus evaluates to

$$- \frac{2}{K} \sum_{k=1}^K \left(\Delta_u^{(k)} - S_u^{(k)} \right) \left(\phi(\xi^{(k)}, \tau^{(k)}) + S_\phi^{(k)} \right). \quad (61)$$

Now, setting $u = V$, Lemma 2.1 readily implies that

$$\frac{\delta}{\delta u} \Big|_{u=V} \widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}(u; \phi) = 0 \quad (62)$$

almost surely for all $\phi \in \mathcal{U}$. This shows that

$$\mathbb{W} \left[\frac{\delta}{\delta u} \Big|_{u=V} \widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}(u; \phi) \right] = 0, \quad (63)$$

which proves the claim. Note that the above calculation shows that the derivative of any squared loss exhibits zero variance whenever the error is vanishing almost surely. \square

Proof of Proposition 2.7. Let us define

$$\delta(x, t) = u(x, t) - V(x, t). \quad (64)$$

Analogously to (61) we can compute

$$\begin{aligned} & \mathbb{W} \left[\frac{\delta}{\delta u} \Big|_{u=V+\delta} \widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}(u; \phi) \right] \\ &= \frac{4}{K} \mathbb{W} [(\delta(\xi, \tau) + S_\delta) (\phi(\xi, \tau) + S_\phi)]. \end{aligned} \quad (65)$$

Now, properties of the variance as well as the Cauchy-Schwarz inequality yield

$$\begin{aligned} & \mathbb{W} [(\delta(\xi, \tau) + S_\delta) (\phi(\xi, \tau) + S_\phi)] \\ & \leq \mathbb{E} \left[(\delta(\xi, \tau) + S_\delta)^2 (\phi(\xi, \tau) + S_\phi)^2 \right] \\ & \leq \mathbb{E} \left[(\delta(\xi, \tau) + S_\delta)^4 \right]^{\frac{1}{2}} \mathbb{E} \left[(\phi(\xi, \tau) + S_\phi)^4 \right]^{\frac{1}{2}}. \end{aligned}$$

Each of the above factors can be bounded by using the Burkholder-Davis-Gundy inequality (Da Prato & Zabczyk, 2014, Section 4.6), Hölder's inequality, and the consistency of the Euclidean norm, i.e.

$$\begin{aligned} & \mathbb{E} \left[(\delta(\xi, \tau) + S_\delta)^4 \right] \leq 8 (\mathbb{E} [\delta(\xi, \tau)^4] + \mathbb{E} [S_\delta^4]) \\ & \leq 8 \left(\varepsilon^4 + 36T \int_0^T \mathbb{E} [\|(\sigma^\top \nabla_x \delta)(X_s, s)\|^4] ds \right) \\ & \leq 8\varepsilon^4 \left(1 + 36T \int_0^T \mathbb{E} [\|\sigma(X_s, s)\|_F^4 (1 + \|X_s\|^\gamma)^4] ds \right). \end{aligned}$$

Defining

$$C := 32 \left(1 + 36T \int_0^T \mathbb{E} [\|\sigma(X_s, s)\|_F^4 (1 + \|X_s\|^\gamma)^4] ds \right),$$

we thus showed that

$$\mathbb{W} \left[\frac{\delta}{\delta u} \Big|_{u=V+\delta} \widehat{\mathcal{L}}_{\text{BSDE}}^{(K)}(u; \phi) \right] \leq \frac{C\varepsilon^2 \kappa^2}{K}, \quad (66)$$

which proves the claim. \square

C. Feynman-Kac Theorem and More General Linear PDEs

As outlined in Section 2 and for instance stated in (5) as well as (13), the Feynman-Kac formula brings a stochastic representation of the linear PDE in (1) via

$$V(x, t) = \mathbb{E}[g(X_T) | X_t = x] \quad (67a)$$

$$= \mathbb{E}[g(X_T) | (\xi, \tau) = (x, t)]. \quad (67b)$$

Note that this can be shown using the identity in (9). More specifically, the stochastic integral S_V has a vanishing expectation conditioned on (ξ, τ) and we obtain that

$$\mathbb{E}[\Delta_V - S_V | (\xi, \tau)] = \mathbb{E}[\Delta_V | (\xi, \tau)] = 0. \quad (68)$$

Let us make this observation precise and at the same time state a slightly more general version of the Feynman-Kac theorem.

Theorem C.1 (Feynman-Kac formula). *Let $g \in C^2(\mathbb{R}^d, \mathbb{R})$, $k \in C(\mathbb{R}^d \times [0, T], \mathbb{R})$, and $V \in C^{2,1}(\mathbb{R}^d \times [0, T], \mathbb{R})$ be at most polynomially growing functions. Further, let*

$f \in C(\mathbb{R}^d \times [0, T], \mathbb{R}_{\geq 0})$ and assume that V solves the parabolic problem

$$(\partial_t + L - f(x, t))V(x, t) + k(x, t) = 0 \quad (69)$$

on $(x, t) \in \mathbb{R}^d \times [0, T]$ with terminal condition

$$V(x, T) = g(x), \quad x \in \mathbb{R}^d. \quad (70)$$

Then

$$V(x, t) = \mathbb{E} \left[\int_t^T e^{-\int_t^r f(X_s, s) ds} k(X_r, r) dr + e^{-\int_t^T f(X_s, s) ds} g(X_T) \middle| X_t = x \right], \quad (71)$$

where X is a strong solution to (6).

Proof. The proof, whose main ingredient is Itô's Lemma, can, for instance, be found in Karatzas & Shreve (1998, Theorem 5.7.6) and Baldi (2017, Theorem 10.5). \square

Let us recall that our loss \mathcal{L}_{FK} as defined in (10) readily follows from the Feynman-Kac formula – see also Beck et al. (2021). Note that a crucial point in the derivation of the Feynman-Kac formula is the elimination of the stochastic integral S_V , as defined in (8b), via its martingale property, see (68). Ironically, as elaborated on in Section 2, it turns out that this elimination is just the reason for larger variances of estimators of \mathcal{L}_{FK} .

Remark C.2 (Initial value problems). Note that time can be reversed and initial value problems (as opposed to terminal value problems as in Theorem C.1) can be formulated. As an example, we can consider the time-homogeneous case, where b, σ, f , and k do not depend on time. Let $V \in C^{2,1}(\mathbb{R}^d \times [0, T], \mathbb{R})$ solve the parabolic PDE

$$(\partial_t - L + f(x))V(x, t) - k(x) = 0 \quad (72)$$

on $(x, t) \in \mathbb{R}^d \times (0, T]$ with initial condition

$$V(x, 0) = g(x), \quad x \in \mathbb{R}^d. \quad (73)$$

The associated stochastic representation is then given by

$$V(x, t) = \mathbb{E} \left[\int_0^t e^{-\int_0^r f(X_s) ds} k(X_r) dr + e^{-\int_0^t f(X_s) ds} g(X_t) \middle| X_0 = x \right]. \quad (74)$$

Remark C.3 (Bounded domains). We can restrict ourselves to open, bounded domains $\mathcal{D} \subset \mathbb{R}^d$ and consider the parabolic PDE in (69) on \mathcal{D} , adding the additional boundary

condition¹⁰ $V(x, t) = g(x)$ for $(x, t) \in \partial\mathcal{D} \times [0, T]$. The stochastic representation then becomes

$$V(x, t) = \mathbb{E} \left[\int_t^{T \wedge \mathcal{T}} e^{-\int_t^r f(X_s, s) ds} k(X_r, r) dr + e^{-\int_t^{T \wedge \mathcal{T}} f(X_s, s) ds} g(X_{T \wedge \mathcal{T}}) \middle| X_t = x \right],$$

where $T \wedge \mathcal{T} := \min\{T, \mathcal{T}\}$ and $\mathcal{T} := \inf\{t \geq 0 : X_t \notin \mathcal{D}\}$ is the first exit time of the stochastic process from the domain \mathcal{D} , for which we usually assume $\mathcal{T} < \infty$ almost surely. Likewise, we can consider the elliptic boundary value problem

$$(L - f(x))V(x) + k(x) = 0, \quad x \in \mathcal{D}, \quad (75a)$$

$$V(x) = g(x), \quad x \in \partial\mathcal{D}, \quad (75b)$$

where now the solution V , the coefficients b and σ in the SDE (6), as well as f and k do not depend explicitly on time anymore, yielding the stochastic representation

$$V(x) = \mathbb{E} \left[\int_0^{\mathcal{T}} e^{-\int_0^r f(X_s) ds} k(X_r) dr + e^{-\int_0^{\mathcal{T}} f(X_s) ds} g(X_{\mathcal{T}}) \middle| X_0 = x \right], \quad (76)$$

again with $\mathcal{T} := \inf\{t \geq 0 : X_t \notin \mathcal{D}\}$, see e.g. Proposition 5.7.2 in Karatzas & Shreve (1998).

Leveraging the above representations, our proposed methods can readily be applied to a range of elliptic and parabolic PDEs on bounded domains. Note, however, that one needs to take into account the hitting times \mathcal{T} , for instance, using naive stopping criteria (Nüsken & Richter, 2021a) or more elaborate walk-on-the-sphere algorithms (Grohs & Herrmann, 2020). As this might obscure the comparisons between different loss functions, we focus on unbounded domains in our experiments.

D. Backward Stochastic Differential Equations and Semi-Linear PDEs

Backward stochastic differential equations (BSDEs) have been studied extensively in the last three decades and we refer to Pardoux (1998), Pham (2009), Gobet (2016), and Zhang (2017) for good introductions to the topic. Even though in this paper we only consider linear PDEs as stated in (1) or (69), BSDEs are typically associated to nonlinear (parabolic) PDEs of the type

$$(\partial_t + L)V(x, t) = h(x, t, V(x, t), (\sigma^\top \nabla_x V)(x, t)) \quad (77)$$

¹⁰One can also consider boundary conditions that are different from the terminal condition in (70), see, for instance, Baldi (2017, Theorem 10.4) and Lelievre & Stoltz (2016, Proposition 6.1).

for $(x, t) \in \mathbb{R}^d \times [0, T]$, a nonlinearity $h : \mathbb{R}^d \times [0, T] \times \mathbb{R} \times \mathbb{R}^d \rightarrow \mathbb{R}$, and differential operator L defined as in (2). The terminal value is given by

$$V(x, T) = g(x), \quad (78)$$

for a specified function $g \in C(\mathbb{R}^d, \mathbb{R})$.

BSDEs were first introduced in [Bismut \(1973\)](#) and their systematic study began with [Pardoux & Peng \(1990\)](#). Loosely speaking, they can be understood as nonlinear extensions of the Feynman-Kac formula (see [Pardoux \(1998\)](#) and [Appendix C](#)), relating the nonlinear PDE in (77) to the stochastic process X_s defined by

$$dX_s = b(X_s, s) ds + \sigma(X_s, s) dW_s, \quad X_0 = x_0. \quad (79)$$

The key idea is then to define the processes

$$Y_s = V(X_s, s), \quad Z_s = (\sigma^\top \nabla_x V)(X_s, s), \quad (80)$$

as representations of the PDE solution and its gradient, respectively, and apply Itô's lemma to obtain

$$dY_s = h(X_s, s, Y_s, Z_s) ds + Z_s \cdot dW_s, \quad (81)$$

with terminal condition $Y_T = g(X_T)$. Noting that the processes Y and Z are adapted¹¹ to the filtration generated by the Brownian motion W , they should indeed be understood as backward processes and not be confused with time-reversed processes. A convenient interpretation of the relations in (80) is that solving for the processes Y and Z under the constraint (81) corresponds to determining the solution of the PDE in (77) (and its gradient) along a random grid which is provided by the stochastic process X defined in (79).

Let us note that under suitable assumptions on the coefficients b, σ, h , and g one can prove existence and uniqueness of a solution to the BSDE system as defined in (79) and (81), see for instance Theorem 4.3.1 in [Zhang \(2017\)](#).

We further note that the standard BSDE system can be generalized to

$$\begin{aligned} dX_s &= (b(X_s, s) + v(X_s, s)) ds + \sigma(X_s, s) dW_s, \\ dY_s &= (h(X_s, s, Y_s, Z_s) + v(X_s, s) \cdot Z_s) ds + Z_s \cdot dW_s, \end{aligned}$$

with

$$X_0 = x, \quad Y_T = g(X_T). \quad (82)$$

In the above, $v \in C(\mathbb{R}^d \times [0, T], \mathbb{R}^d)$ is a suitable control vector field that can be understood as pushing the forward trajectories into desired regions of the state space, noting that the relations

$$Y_s = V(X_s, s), \quad Z_s = (\sigma^\top \nabla_x V)(X_s, s), \quad (83)$$

¹¹Intuitively, this means that the processes Y and Z must not depend on future values of the Brownian motion W .

Algorithm 1 Solving the PDE in (1) via deep learning

input Neural network Φ with initial parameters $\theta^{(0)}$, optimizer method step for updating the parameters, maximum number of steps M , batch size K , step-size Δt
output parameters $\theta^{(M)}$

for $m \leftarrow 0, \dots, M - 1$ **do**
 $(\xi^{(k)}, \tau^{(k)}, W^{(k)})_{k=1}^K \leftarrow$ sample from $(\xi, \tau, W)^{\otimes K}$
 $(\widehat{X}^{(k)})_{k=1}^K \leftarrow$ simulate using the EM scheme in (84)
 $\mathcal{L} \leftarrow$ pick $\mathcal{L} \in \{\mathcal{L}_{\text{FK}}, \mathcal{L}_{\text{BSDE}}\}$
 $\widehat{\mathcal{L}}^{(K)}(\Phi_{\theta^{(m)}}) \leftarrow$ compute estimator loss as in (15)
 $\nabla_{\theta} \widehat{\mathcal{L}}^{(K)}(\Phi_{\theta^{(m)}}) \leftarrow$ autodiff($\widehat{\mathcal{L}}^{(K)}(\Phi_{\theta^{(m)}})$)
 $\theta^{(m+1)} \leftarrow$ step($\theta^{(m)}, \nabla_{\theta} \widehat{\mathcal{L}}^{(K)}(\Phi_{\theta^{(m)}})$)

end for

with $V \in C^{2,1}(\mathbb{R}^d \times [0, T], \mathbb{R})$ being the solution to the parabolic PDE in (77), hold true independent of the choice of v ([Hartmann et al., 2019](#)).

E. Further Computational Details

For convenience, we first summarize our considered method in [Algorithm 1](#). In numerical simulations we need to discretize the stochastic process X as defined in (6) on a time grid $\tau = t_1 < \dots < t_J$. A practical way to do so is based on the Euler-Maruyama (EM) scheme

$$\widehat{X}_{j+1} = \widehat{X}_j + b(\widehat{X}_j, t_j) \Delta t + \sigma(\widehat{X}_j, t_j) \sqrt{\Delta t} \zeta_{j+1}, \quad (84)$$

where $\Delta t := t_{j+1} - t_j$ is the step-size and

$$\zeta_{j+1} := \frac{W_{t_{j+1}} - W_{t_j}}{\sqrt{\Delta t}} \sim \mathcal{N}(0, \text{Id}) \quad (85)$$

is a standard normally distributed random variable. It can be shown that \widehat{X}_j converges to $X_{j\Delta t}$ in an appropriate sense ([Kloeden & Platen, 1992](#)). This readily leads to discrete versions of the quantities Δ_u and S_u as defined in (16). Note that the discrete process \widehat{X} is initialized at the random value $\widehat{X}_1 = \xi$ and J is chosen according to the randomly drawn initial time τ . More precisely, we set $J = \lceil (T - \tau) / \Delta t \rceil$ and use a smaller final step-size in order to arrive at the terminal time T . An alternative strategy would be to fix J for all realizations and change the step-size $\Delta t = \frac{T - \tau}{J}$ depending on the value of τ . We display the memory requirements of our methods as a function of the step-size Δt in [Figure 7](#).

In [Table 1](#) we summarize the hyperparameters for our experiments. We specify the amount of steps ([Schedule A](#)) or time ([Schedule B](#)) that we train and the percentile, i.e., milestone, where we decrease the learning rate λ and the step-size Δt . In order to have comparable results, we set a GPU memory limit of 8 GiB during training and always use batch sizes $K \in \{2^4, 2^7, 2^{10}, 2^{13}, 2^{16}, 2^{19}\}$.




Figure 7. GPU memory requirements for a gradient step with batch size $K = 1024$ and different step-sizes Δt for the SDE discretization when solving the heat equation. If we do not consider the efficient versions, the memory usage for the losses $\mathcal{L}_{\text{BSDE}}^{\text{grad}}$ and $\mathcal{L}_{\text{BSDE}}$ depends approximately linearly on Δt .

We employ a Multilevel neural network architecture, which has shown to be advantageous over standard feed-forward architectures. More precisely, we use an architecture with $L = 3$ levels, amplifying factor $q = 5$ for the HJB equation and $q = 3$ for the other PDEs, intermediate residual connections ($\chi = 1$) and without normalization layer, as defined in Berner et al. (2020a). Note that in case of the losses $\mathcal{L}_{\text{BSDE}}^{\text{grad}}$ and $\mathcal{L}_{\text{BSDE},\text{eff}}^{\text{grad}}$ we use the same architecture with output dimension d for the neural network representing the function r in (31). As the activation function needs to be twice differentiable for certain losses, we replace the ReLU by the SiLU (also known as swish) activation function (Hendrycks & Gimpel, 2016).

Finally, Table 1 also specifies the number of samples used to approximate the MSE metrics and the number of batches used to estimate the loss and gradient variances. For the latter, we compute the value and derivative of $\hat{\mathcal{L}}^{(K)}$ (w.r.t. the parameters θ of Φ_θ) for 30 batches consisting of K independent samples of (ξ, τ, W) by performing forward and backward passes without updating the parameters θ . To compute the MSE we evaluate $\hat{\mathcal{L}}_{\text{Eval}}$ with $10 \cdot 2^{17}$ i.i.d. samples drawn from the distribution of (ξ, τ) – independently of the training data and independently for each evaluation. In this sense, we always evaluate our model w.r.t. to the real solution V on unseen data. For comparing the gradients we replace $(u(\xi, \tau) - V(\xi, \tau))^2$ by $\|\nabla_x u(\xi, \tau) - \nabla_x V(\xi, \tau)\|^2$ or $\|r(\xi, \tau) - \nabla_x V(\xi, \tau)\|^2$ in case of the methods $\mathcal{L}_{\text{BSDE}}^{\text{grad}}$ and $\mathcal{L}_{\text{BSDE},\text{eff}}^{\text{grad}}$, respectively. When no closed-form solution for V is available, as for the Black–Scholes model in Section 3.2, we use the version in (7b) with another 2^{16} independent samples to estimate the inner expectation.

F. Further Numerical Experiments

Figures 8 to 15 show the results of additional numerical experiments. The corresponding captions describe the respective settings.

Table 1. Training and evaluation setup

Schedule A: Step limit	
steps	$3 \cdot 10^4$
milestone	0.9
Schedule B: Time limit	
time	24h
milestone	0.5
Training	
learning rate λ	$[5 \cdot 10^{-4}, 5 \cdot 10^{-6}]$
step-size Δt	$[10^{-2}, 10^{-3}]$
optimizer	Adam
batch size K	$\{2^4, 2^7, 2^{10}, 2^{13}, 2^{16}, 2^{19}\}$
GPU memory constraint	8 GiB
Network	
architecture	Multilevel
(L, q, χ)	$\{(3, 3, 1), (3, 5, 1)\}$
activation function	SiLU
Evaluation	
samples	$10 \cdot 2^{17}$
batches (for variances)	30




Figure 8. Approximation of the norm of the gradient $\|\nabla_x V\|$ of the solution V to the heat equation in Section 3.1 for $t \in \{0.25, 0.75\}$ and $x = (\alpha, \dots, \alpha)^\top$ after training for $30k$ steps with batch size $K = 1024$. As expected, the stochastic integral, which is present in all losses except \mathcal{L}_{FK} , induces better approximation of the gradient. Furthermore, explicitly modelling the gradient as in $\mathcal{L}_{\text{BSDE},\text{eff}}^{\text{grad}}$ or back-propagating the derivative of the neural network as in $\mathcal{L}_{\text{BSDE},\text{eff}}$ also improves the approximation outside of the sampling interval for x , i.e. outside of $[-0.5, 0.5]^{50}$.




Figure 9. Estimator standard deviation and MSE after $30k$ gradient steps for the heat equation in Section 3.1 with batch size 1024 and varying diffusivities ν . While solving the PDE becomes more challenging for higher values of ν , our proposed methods consistently outperform the baseline.




Figure 10. We observe performance differences similar to Figure 1 when solving the heat equation from Section 3.1 with a higher diffusivity of $\nu = \sqrt{3}$.




Figure 11. MSE of different losses as a function of the training time when solving the heat equation from Section 3.1 with batch size $K = 8192$. Although the time needed for computing and back-propagating the stochastic integral leads to significantly less gradient steps and thus samples of (ξ, τ, W) (see Figure 12), the loss $\mathcal{L}_{\text{BSDE, eff}}$ still outperforms the loss \mathcal{L}_{FK} .

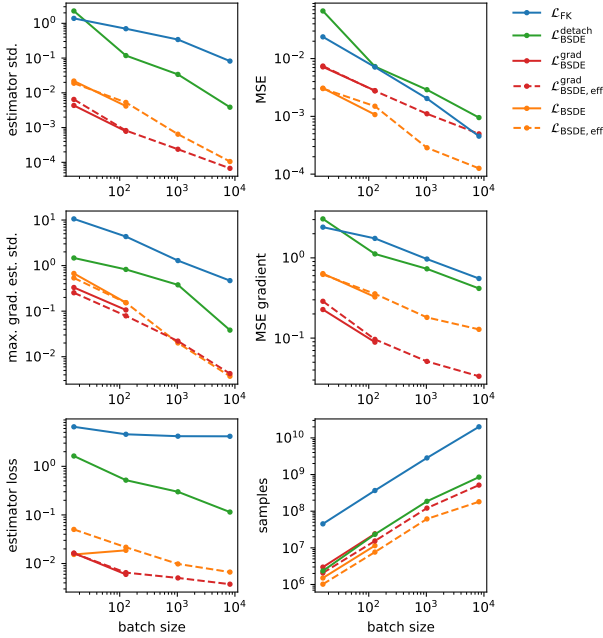


Figure 12. Scaling of estimator loss, standard deviation of the (gradient) estimator, and MSE when training for 24h. We observe similar effects as in Figures 1 and 4, even though the number of samples (or, proportionally, the number of steps) is significantly higher for the loss \mathcal{L}_{FK} . For batch sizes higher than $K = 8192$ a larger time budget is necessary.




Figure 13. Performance and standard deviation of the losses $\mathcal{L}_{\text{Eval}}$ and $\mathcal{L}_{\text{Eval}}^{\text{detach}}$ (detaching S_v in (7c) from the computational graph as in (30)) for the heat equation from Section 3.1 compared to their natural counterparts. Note that they perform significantly worse for small batch sizes, but similarly for large batch sizes – for the detached version one can allow for larger batch sizes, given a fixed memory budget. Note via a comparison to Figures 1 and 4 that the efficient versions of the considered losses are still much better.




Figure 14. Estimator standard deviation and MSE as a function of the gradient steps when solving the Black-Scholes equation from Section 3.2 with batch size $K = 8192$. In case of the loss $\mathcal{L}_{\text{BSDE}^{\text{detach}}}$, a bad initial approximation actually leads to a variance-increasing effect of the stochastic integral. As the error in the gradient is not back-propagated to the network parameters, this deteriorates the convergence.




Figure 15. Similar to Figures 1, 4, and 5, our theoretical and empirical findings also hold in case of the HJB equation from Section 3.3.

G. Additional Material

The following Lemma details the relation of the HJB equation from Section 3.3 and a linear PDE of Kolmogorov type.

Lemma G.1 (Linearization of HJB equation). *Let $\tilde{V} \in C^{2,1}(\mathbb{R}^d \times [0, T], \mathbb{R})$ solve the Hamilton–Jacobi–Bellman equation in (45), then $V = \exp(-\tilde{V})$ fulfills the linear PDE as stated in Theorem C.1 with $k = 0$ and $g = \exp(-\tilde{g})$, i.e.*

$$\begin{aligned} (\partial_t + L - f(x, t))V(x, t) &= 0, & (x, t) &\in \mathbb{R}^d \times [0, T], \\ V(x, T) &= g(x), & x &\in \mathbb{R}^d. \end{aligned}$$

Subsequently, we get the representation

$$\tilde{V}(x, t) = -\log \mathbb{E} \left[e^{-\mathcal{W}(X)} \middle| X_t = x \right] \quad (87)$$

with $\mathcal{W}(X) = \int_t^T f(X_s, s) ds + \tilde{g}(X_T)$.

Proof. We consider the transformation $V = \exp(-\tilde{V})$ and for notational convenience omit the space and time dependencies of \tilde{V} , f , b , and σ . We can compute

$$\begin{aligned} Le^{-\tilde{V}} &= -b \cdot e^{-\tilde{V}} \nabla_x \tilde{V} - \frac{1}{2} \sum_{i,j=1}^d (\sigma \sigma^\top)_{ij} \partial_{x_i} (\partial_{x_j} e^{-\tilde{V}} \tilde{V}) \\ &= -e^{-\tilde{V}} \left(b \cdot \nabla_x \tilde{V} + \frac{1}{2} \sum_{i,j=1}^d (\sigma \sigma^\top)_{ij} \partial_{x_i} \partial_{x_j} \tilde{V} \right. \\ &\quad \left. - \frac{1}{2} \sum_{i,j=1}^d (\sigma \sigma^\top)_{ij} \partial_{x_i} \tilde{V} \partial_{x_j} \tilde{V} \right) \\ &= -e^{-V} \left(b \cdot \nabla_x \tilde{V} + \frac{1}{2} \sum_{i,j=1}^d (\sigma \sigma^\top)_{ij} \partial_{x_i} \partial_{x_j} \tilde{V} \right. \\ &\quad \left. - \frac{1}{2} \sum_{i,j,k=1}^d \sigma_{ik} \sigma_{jk} \partial_{x_i} \tilde{V} \partial_{x_j} \tilde{V} \right) \\ &= -e^{-\tilde{V}} \left(b \cdot \nabla_x \tilde{V} + \frac{1}{2} \sum_{i,j=1}^d (\sigma \sigma^\top)_{ij} \partial_{x_i} \partial_{x_j} \tilde{V} \right. \\ &\quad \left. - \frac{1}{2} \sum_{k=1}^d \left(\sum_{i=1}^d \sigma_{ik} \partial_{x_i} \tilde{V} \right)^2 \right) \\ &= -e^{-\tilde{V}} \left(L\tilde{V} - \frac{1}{2} \|\sigma^\top \nabla_x \tilde{V}\|^2 \right). \end{aligned}$$

The PDE in (69) therefore becomes

$$0 = (\partial_t + L - f)e^{-\tilde{V}} \quad (88a)$$

$$= -e^{-\tilde{V}} \left((\partial_t + L)\tilde{V} + f - \frac{1}{2} \|\sigma^\top \nabla_x \tilde{V}\|^2 \right), \quad (88b)$$

which is equivalent to the HJB equation in (45). \square

For a discussion on the application of Lemma G.1 in the context of importance sampling of stochastic processes that are relevant in computational statistical physics and molecular dynamics we refer to Hartmann et al. (2017).