A NEURAL NETWORK APPROACH FOR STOCHASTIC OPTIMAL CONTROL*

XINGJIAN LI†, DEEPANSHU VERMA†, AND LARS RUTHOTTO†

Abstract. We present a neural network approach for approximating the value function of high-dimensional stochastic control problems. Our training process simultaneously updates our value function estimate and identifies the part of the state space likely to be visited by optimal trajectories. Our approach leverages insights from optimal control theory and the fundamental relation between semi-linear parabolic partial differential equations and forward-backward stochastic differential equations. To focus the sampling on relevant states during neural network training, we use the stochastic Pontryagin maximum principle (PMP) to obtain the optimal controls for the current value function estimate. By design, our approach coincides with the method of characteristics for the non-viscous Hamilton-Jacobi-Bellman equation as the dynamics become deterministic. Our training loss consists of a weighted sum of the objective functional of the control problem and penalty terms that enforce the HJB equations along the sampled trajectories. Importantly, training is unsupervised in that it does not require solutions of the control problem.

In our numerical experiments, we compare our method to existing solvers for a more general class of semi-linear PDEs. Using a two-dimensional model problem, we demonstrate the importance of the stochastic PMP to inform the sampling. For a 100-dimensional benchmark problem, we demonstrate that our approach improves accuracy and time-to-solution.

Key words. Hamilton-Jacobi-Bellman equation, high-dimensional stochastic optimal control, forward-backward stochastic differential equations, neural networks, stochastic maximum principle

AMS subject classifications. 35F21, 49M99, 68T07

1. Introduction. We consider Stochastic Optimal Control (SOC) problems that require finding a policy to control randomly perturbed dynamical systems so as to optimize a given objective functional. Problems of this type arise in many fields including finance, biology, robotics and many other engineering applications; see, e.g., [10, 34] for references and extensive theoretical discussion on SOC problems.

Dynamic programming (DP) is a prominent framework for solving SOC problems. At its core, DP seeks to find the value function, which assigns every state of the system the optimal cost-to-go. For many problems, recovering the optimal control from the value function is straightforward. One way to compute the value function is by solving the Hamilton-Jacobi-Bellman (HJB) equation [2, 37], which is a semi-linear parabolic Partial Differential Equation (PDE). The two main challenges in solving the HJB equations is the forward-backward structure and the high dimensionality.

The space dimension of the HJB equation equals the state dimension of the dynamical system to be controlled. Hence, many numerical schemes for solving PDEs cannot be applied to realistic problem instances due to the Curse-of-Dimensionality (CoD); for example, the computational costs of approaches that discretize the state space with a mesh typically grows exponentially in the dimensions of the system. Hence, the applicability of existing approaches that employ spatial discretizations (for example, [4, 6, 16, 20, 32, 22, 23, 32, 35]) is limited to state dimensions $d \leq 3$. Mesh-free methods based on radial basis function methods (for example, [36]) can be effective for slightly larger d but ultimately also suffer from CoD for $d = \mathcal{O}(100)$, as

†Department of Mathematics, Emory University, Atlanta, GA (xingjian.li@emory.edu, deepanshu.verma@emory.edu,lruthotto@emory.edu).

we consider here.

A first requirement for mitigating the CoD in SOC problems is to alleviate the need for a spatial discretization. This can be achieved by using a nonlinear version of the Feynman-Kac lemma and replacing the HJB equation by a system of Forward-Backward Stochastic Differential Equations (FBSDEs); see, for example, [5, 28, 37]. Discretizing the FBSDE system using the classical Euler Maruyama scheme [19] yields mesh-free numerical schemes that have gained importance for SOC problems and non-linear second-order PDEs; see, for example, [3, 8, 9, 13, 25] and references therein.

A second requirement for mitigating the CoD is to parameterize the value function effectively in high-dimensions. Due to their universal approximation property and many advances in deep learning, the use of neural networks (NNs) as function approximators has been gathering significant attention recently. The idea of combining the BSDE approach with NNs has been pioneered by the seminal works [7, 12] that use neural networks to solve high-dimensional semi-linear parabolic PDEs and include an example for HJB equations as a special case. Around the same time, [33] achieves similarly promising results for neural-network based solution of the HJB equation. These works developed learning algorithms for approximating the value function or its gradient using the FBSDE system. Their success has sparked several extensions that improve the NN architectures and consider other types of PDEs; see, for example, [15, 17, 29].

Even after combining FBSDEs and deep learning, solving the HJB equation globally remains affected by the curse of dimensionality as it would require sampling the entire state-space. Fortunately, the optimal state trajectories of many SOC problems do not cover the entire space. This motivates us to develop semi-global approaches, that is, we seek to reliably estimate the value function in those parts of the state space that are likely to be visited when following an optimal policy. Clearly, these tasks are interrelated: Finding the relevant parts of the state space depends on the value function and, vice versa, the value function needs to be trained using samples from the state space.

The key idea of our approach is to use the stochastic Pontryagin maximum principle (PMP) [31] to define the forward system in the FBSDE approach in terms of the value function. The PMP provides a set of necessary optimality conditions, known as the Hamiltonian system, and a feedback form that allows recovering the optimal control from the value function; see, for example, [37]. As we discuss in more detail later, the choice of the forward SDE is crucial as it is used to sample the parts of the state space along which the value function estimate is improved using a loss that includes the backward SDE. It is worth noting that the FBSDE framework allows choosing the forward system almost arbitrarily. Thus, our choice is theoretically on par with the standard Brownian motion used for exploring the state space in [12, 33], but as our numerical experiments suggest enable solving a wider class of practical problems; see section 4.

A theoretical advantage of our proposed scheme over existing approaches is its consistency with deterministic optimal control problems. In the absence of uncertainty in the dynamics, our approach reduces to the method of characteristics for the HJB equation. Hence our approach can also be seen as an extension of the neural network approaches for deterministic optimal control proposed in [27, 21].

The rest of the paper is organised as follows: In section 2 we introduce the parts of SOC theory used to obtain our proposed forward SDE. In section 3, we provide the NN approach to learn the value function using the FBSDE reformulation from section 2. Using an intuitive 2D example, we show the importance of sampling. To

illustrate the potential of our method, we consider the 100-dimensional benchmark problem also used in [12, 33], propose a modification that leads to a more challenging sampling problem, and investigate the scaling of our approach with respect to problem dimension. Finally we conclude the paper and discuss future directions.

2. Stochastic Optimal Control Background. In this section, we describe the Stochastic Optimal Control (SOC) problems considered in this work and review the key results from SOC theory that motivates our approach; our discussion follows [37] and we refer to this text book for a more comprehensive background. We first introduce the SOC problem and then review its underlying theory; to be specific, we review the stochastic Pontryagin Maximum Principle (PMP), the Hamilton-Jacobi-Bellman (HJB) equation, and its reformulation into a system of forward backward stochastic differential equations (FBSDEs) obtained from a nonlinear version of the Feynman-Kac formula.

2.1. Stochastic Optimal Control Problem. Let $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P})$ be a given complete probability space, $W(s)$ be a d-dimensional Brownian motion on $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ where s denotes the time. For a fixed initial state $\boldsymbol{x}$ at some time $0 < t < T < \infty$, we seek to control the randomly perturbed dynamical system

(2.1)
$$\begin{cases} \mathrm{d}\boldsymbol{z}_{t,\boldsymbol{x}}(s) = f(s, \boldsymbol{z}_{t,\boldsymbol{x}}(s), \boldsymbol{u}_{t,\boldsymbol{x}}(s, \boldsymbol{z}_{t,\boldsymbol{x}}(s)))\mathrm{d}s + \sigma(s, \boldsymbol{z}_{t,\boldsymbol{x}}(s))\mathrm{d}W(s), \quad t \leq s \leq T, \\ \boldsymbol{z}_{t,\boldsymbol{x}}(t) = \boldsymbol{x}. \end{cases}$$

Here, $\boldsymbol{z}_{t,\boldsymbol{x}} : [t, T] \to \mathbb{R}^d$ describes the state and $\boldsymbol{u}_{t,\boldsymbol{x}} : [t, T] \times \mathbb{R}^d \to U$ describes the control of the system, the function $\sigma : [t, T] \times \mathbb{R}^d \to \mathbb{R}^{d\times d}$ represents the diffusion coefficient, and $f : [t, T] \times \mathbb{R}^d \times U \to \mathbb{R}^d$ represents the drift of the system. We assume that the set of admissible controls $U \subset \mathbb{R}^k$ is closed. We seek to minimize the objective functional

(2.2) $$J_{t,\boldsymbol{x}}(\boldsymbol{u}_{t,\boldsymbol{x}}) = \mathbb{E}\left\{ G(\boldsymbol{z}_{t,\boldsymbol{x}}(T)) + \int_t^T L(s, \boldsymbol{z}_{t,\boldsymbol{x}}(s), \boldsymbol{u}_{t,\boldsymbol{x}}(s, \boldsymbol{z}_{t,\boldsymbol{x}}(s)))\,\mathrm{d}s \right\},$$

which is comprised of the running cost $L : [t, T] \times \mathbb{R}^d \times U \to \mathbb{R}$ and the terminal cost $G : \mathbb{R}^d \to \mathbb{R}$. Here, the expectation is taken with respect to perturbances of the dynamics (2.1) that is described by the probability measure $\mathbb{P}$. We assume sufficient regularity conditions on f, σ, G and L, see [37, Chapter 2] for a list of assumptions.

The value function assigns the optimal cost-to-go to any initial state, that is,

(2.3) $$\Phi(t, \boldsymbol{x}) = \inf_{\boldsymbol{u}_{t,\boldsymbol{x}}} J_{t,\boldsymbol{x}}(\boldsymbol{u}_{t,\boldsymbol{x}}),$$

and a solution $\boldsymbol{u}_{t,\boldsymbol{x}}^*$ to (2.3) incurring this minimum value is called an optimal control.

The *generalized Hamiltonian*, $H : [t, T] \times \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^{d\times d} \to \mathbb{R} \cup \{\infty\}$, is a key ingredient for the SOC theory in the following sections and the backbone of our numerical scheme. For the problem defined in (2.1) and (2.2) it reads

(2.4) $$H(s, \boldsymbol{z}, \boldsymbol{p}, \boldsymbol{M}) = \sup_{\boldsymbol{u}\in U} \mathcal{H}(s, \boldsymbol{z}, \boldsymbol{p}, \boldsymbol{M}, \boldsymbol{u}),$$

where $\boldsymbol{p}$ and $\boldsymbol{M}$ are called adjoint variables and

$$\mathcal{H}(s, \boldsymbol{z}, \boldsymbol{p}, \boldsymbol{M}, \boldsymbol{u}) = \frac{1}{2}\mathrm{tr}\left(\sigma(s, \boldsymbol{z})\boldsymbol{M}\right) + \boldsymbol{p}\cdot f(s, \boldsymbol{z}, \boldsymbol{u}) - L(s, \boldsymbol{z}, \boldsymbol{u}).$$

To make it notationally convenient, in the rest of the paper, we drop the second argument for the controls and denote controls by $\boldsymbol{u}_{t,\boldsymbol{x}}(s)$.

2.2. Stochastic Pontryagin Maximum Principle. The PMP provides first-order necessary conditions for the SOC problem and also states that the optimal control $\boldsymbol{u}^*_{t,\boldsymbol{x}}$ must satisfy an (extended) Hamiltonian system along the optimal state and adjoint trajectory. This is made precise by the following result from [37, Theorem 3.2, Chapter 3].

THEOREM 2.1. *[37, Theorem 3.2, Chapter 3] Assume that $(\boldsymbol{z}^*_{t,\boldsymbol{x}}, \boldsymbol{u}^*_{t,\boldsymbol{x}})$ is an optimal pair that solves (2.1) and (2.2). Then there exist adjoint states $\boldsymbol{p}_{t,\boldsymbol{x}} \colon [t,T] \to \mathbb{R}^d$ and $\boldsymbol{M}_{t,\boldsymbol{x}} \colon [t,T] \to \mathbb{R}^{d \times d}$ satisfying the adjoint equation*

$$
(2.5) \quad \begin{cases} \mathrm{d}\boldsymbol{p}_{t,\boldsymbol{x}}(s) = & \boldsymbol{M}_{t,\boldsymbol{x}}(s)\mathrm{d}W(s) - \nabla_{\boldsymbol{z}}\mathcal{H}\big(s, \boldsymbol{z}^*_{t,\boldsymbol{x}}(s), \boldsymbol{p}_{t,\boldsymbol{x}}(s), \boldsymbol{M}_{t,\boldsymbol{x}}(s), \boldsymbol{u}^*_{t,\boldsymbol{x}}(s)\big)\mathrm{d}s \\ \boldsymbol{p}_{t,\boldsymbol{x}}(T) = & -\nabla_{\boldsymbol{z}}G\big(\boldsymbol{z}^*_{t,\boldsymbol{x}}(T)\big), \end{cases}
$$

where $s \in [t,T]$ and the optimal control satisfies

$$
(2.6) \quad \boldsymbol{u}^*_{t,\boldsymbol{x}}(s) \in \underset{\boldsymbol{u}\in U}{\arg\max} \ \mathcal{H}\big(s, \boldsymbol{z}^*_{t,\boldsymbol{x}}(s), \boldsymbol{p}_{t,\boldsymbol{x}}(s), \boldsymbol{M}_{t,\boldsymbol{x}}(s), \boldsymbol{u}(s)\big)
$$

for almost all $s \in [t,T]$, $\mathbb{P}$-almost surely.

We note that the optimal control defined in (2.6) only depends on the adjoint variable $\boldsymbol{p}_{t,\boldsymbol{x}}$ but not on $\boldsymbol{M}_{t,\boldsymbol{x}}$ since the diffusion $\sigma(\cdot, \cdot)$ does not depend on the control in our case.

We further assume that there exists a unique continuous closed form solution to (2.6). Although not demonstrated in this work, this assumption can be weakened to include implicitly defined functions as long as they can be obtained efficiently; this allows, for example, modeling more general convex running costs.

We note that when the control satisfies (2.6), the dynamics in (2.1) is equal to

$$
(2.7)
$$
$$
\begin{cases} \mathrm{d}\boldsymbol{z}^*_{t,\boldsymbol{x}}(s) = \nabla_{\boldsymbol{p}}\mathcal{H}\big(s, \boldsymbol{z}^*_{t,\boldsymbol{x}}(s), \boldsymbol{p}_{t,\boldsymbol{x}}(s), \boldsymbol{M}_{t,\boldsymbol{x}}(s), \boldsymbol{u}^*_{t,\boldsymbol{x}}(s)\big)\mathrm{d}s + \sigma(s, \boldsymbol{z}^*_{t,\boldsymbol{x}}(s))\mathrm{d}W(s), \\ \boldsymbol{z}^*_{t,\boldsymbol{x}}(t) = \boldsymbol{x}. \end{cases}
$$

The system of equations (2.5)–(2.7) is called the stochastic Hamiltonian system, where the maximum condition (2.6) corresponds to the variational inequality for the control.

Finding a tuple $(\boldsymbol{z}^*_{t,\boldsymbol{x}}, \boldsymbol{u}^*_{t,\boldsymbol{x}}, \boldsymbol{p}_{t,\boldsymbol{x}}, \boldsymbol{M}_{t,\boldsymbol{x}})$ that satisfies the PMP can be extremely difficult. However, when the value function Φ is differentiable, $(\boldsymbol{p}_{t,\boldsymbol{x}}, \boldsymbol{M}_{t,\boldsymbol{x}})$ satisfying (2.5) can be obtained from Φ; this is formalized in the following theorem that, with weaker assumptions, can be found in [37, Chapter 5].

THEOREM 2.2. *[37, Chapter 5] Assume that $\boldsymbol{u}^*_{t,\boldsymbol{x}}$ is an optimal control and $\Phi \in C^{1,3}([t,T] \times \mathbb{R}^d)$. Then*
$$
(2.8)
$$
$$
\boldsymbol{p}_{t,\boldsymbol{x}}(s) = -\nabla_{\boldsymbol{z}}\Phi\big(s, \boldsymbol{z}^*_{t,\boldsymbol{x}}(s)\big) \quad and \quad \boldsymbol{M}_{t,\boldsymbol{x}}(s)) = -\sigma(s, \boldsymbol{z}^*_{t,\boldsymbol{x}}(s))^{\top}\nabla^2_{\boldsymbol{z}}\Phi\big(s, \boldsymbol{z}^*_{t,\boldsymbol{x}}(s)\big)
$$

solve (2.5).

Theorem 2.2 simplifies (2.6) to

$$
(2.9) \quad \boldsymbol{u}^*_{t,\boldsymbol{x}}(s) = \boldsymbol{u}^*_{t,\boldsymbol{x}}\left(s, \boldsymbol{z}^*_{t,\boldsymbol{x}}(s), -\nabla_{\boldsymbol{z}}\Phi\big(s, \boldsymbol{z}^*_{t,\boldsymbol{x}}(s)\big)\right).
$$

This relation along with (2.7) is one of the key ingredients of our numerical solution approach. Equation (2.9) characterizes optimal control in a feedback or closed-loop

form, which is of utmost importance in many real-life applications. This is because optimal control can be quickly computed at any given point in time and space, when the value function is known and its gradient is readily available. Thereby, avoiding re-computation of optimal controls for cases when multiple evaluations are needed for different time or states.

2.3. Hamilton-Jacobi-Bellman Equation. To help approximate the value function Φ, we also use the fact that Φ satisfies the Hamilton-Jacobi-Bellman (HJB) equation, which is a result of the Dynamic Programming (DP) method or Bellman's principle. We state the following result taken from [37] under suitable assumptions, see also [30, Remark 3.4.4, Theorem 3.5.2].

THEOREM 2.3. *[37, Propositon 3.5, Chapter 4] Assume that the value function* $\Phi \in C^{1,2}([t,T] \times \mathbb{R}^d)$. *Then Φ satisfies the HJB equation*

(2.10)
$$-\partial_s \Phi(s,\boldsymbol{z}) + H\big(s,\boldsymbol{x}, -\nabla_{\boldsymbol{z}}\Phi(s,\boldsymbol{z}), -\sigma(s,\boldsymbol{z})^\top \nabla_{\boldsymbol{z}}^2 \Phi(s,\boldsymbol{z})\big) = 0, \ \forall \, (s,\boldsymbol{z}) \in [t,T) \times \mathbb{R}^d,$$
$$\Phi(T,\boldsymbol{z}) = G(\boldsymbol{z}).$$

The smoothness of Φ can be relaxed to continuity in the weaker sense of viscosity solutions [37, Section 5, Chapter 4].

Using the relation between PMP and DP from Theorem 2.2 the HJB equation, (2.10), can now be re-written as the following second order parabolic PDE:

(2.11)
$$\begin{cases} -\partial_s \Phi(s,\boldsymbol{z}) - \dfrac{1}{2}\mathrm{tr}(\sigma(s,\boldsymbol{z})\sigma(s,\boldsymbol{z})^\top \nabla^2 \Phi(s,\boldsymbol{z})) - \nabla\Phi(s,\boldsymbol{z}) \cdot f(s,\boldsymbol{z},\boldsymbol{u}^*) \\ \qquad\qquad\qquad\qquad\qquad\qquad - L(s,\boldsymbol{z},\boldsymbol{u}^*) = 0, \ \forall \, (s,\boldsymbol{z}) \in [t,T) \times \mathbb{R}^d, \\ \Phi(T,\boldsymbol{z}) = G(\boldsymbol{z}(T)). \end{cases}$$

In addition, by the envelope theorem, it follows that $\nabla_{\boldsymbol{p}}\mathcal{H} = \nabla_{\boldsymbol{p}}H$ and $\nabla_{\boldsymbol{M}}\mathcal{H} = \nabla_{\boldsymbol{M}}H$. This simplifies the computation of optimal trajectories, which can now be expressed via the value function as:

(2.12)
$$\begin{cases} \mathrm{d}\boldsymbol{z}_{t,\boldsymbol{x}}^*(s) = \nabla_{\boldsymbol{p}}H\big(s, \boldsymbol{z}_{t,\boldsymbol{x}}^*(s), -\nabla_{\boldsymbol{z}}\Phi\big(s, \boldsymbol{z}_{t,\boldsymbol{x}}^*(s)\big), -\sigma(s, \boldsymbol{z}_{t,\boldsymbol{x}}^*(s))^\top \nabla_{\boldsymbol{z}}^2 \Phi\big(s, \boldsymbol{z}_{t,\boldsymbol{x}}^*(s)\big)\big)\mathrm{d}s \\ \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad + \sigma(s, \boldsymbol{z}_{t,\boldsymbol{x}}^*(s))\mathrm{d}W(s), \\ \boldsymbol{z}_{t,\boldsymbol{x}}^*(t) = \boldsymbol{x}. \end{cases}$$

These modified dynamics do not explicitly involve the control, which reduces the problem solely to state variable. This idea was proposed for deterministic optimal control problems in [27].

2.4. FBSDE Formulation. One way to avoid the need for a spatial discretization of the HJB equation (2.11) is to use a non-linear version of Feynman-Kac formula and obtain an equivalent system of stochastic differential equations. This idea has been applied to a variety of nonlinear parabolic/elliptic PDEs; see, e.g., [1, 24, 5, 28, 30].

Our FBSDE system uses (2.12) as the forward system to sample trajectories. Along those trajectories, we note that the solution to the HJB equation (2.11) must

satisfy the backward SDE

$$(2.13) \quad \begin{cases} \mathrm{d}\Phi(s, \boldsymbol{z}(s)) &= \nabla_{\boldsymbol{z}}\Phi(s, \boldsymbol{z}(s))^{\top}\sigma(s, \boldsymbol{z}(s))\,\mathrm{d}W(s) - L(s, \boldsymbol{z}(s), \boldsymbol{u}^*(s))\mathrm{d}s, \\ \Phi(T, \boldsymbol{z}(T)) &= G(\boldsymbol{z}(T)). \end{cases}$$

It is important to stress that our choice of the forward system (2.12) is the key difference to other existing works. For example, [12] and [33] use the standard Brownian motion. While both of these choices lead to a valid FBSDE system for (2.11), we advocate for including the control in the dynamics as motivated by stochastic PMP (2.7). As our numerical experiments demonstrate, focusing the sampling along optimal trajectories can lead to more accurate and efficient value function approximations.

3. Neural Network Approach. In this section, we present a neural network framework for approximating the value function of the stochastic optimal control problem defined by the objective functional (2.2) and dynamics (2.1). The theoretical foundation of our framework are given by the PMP, FBSDE, and Dynamic Programming as presented in the previous section. The key idea is to approximate the value function Φ in (2.3) by a neural network and compute the control using the feedback form (2.9). What sets our framework apart is the use of the feedback form to guide the sampling during training. Thereby we seek to learn to explore the relevant part of the state space. We also derive and experiment with various loss functions that are based on the control objective (2.2), the BSDE (2.13), and the HJB (2.11).

3.1. Neural Network Approximation. The first building block of our framework is to parameterize the value function using a neural network. Since finding an effective network architecture for any learning task is both crucial and an open research topic, we treat this as a modular component. Our framework can be used with any scalar-valued neural network that takes inputs in $\mathbb{R}^{d+1}$ as long as it is twice differentiable with respect to its inputs.

Among the networks we use in our numerical experiments is the Multi-Layer Perceptron (MLP) model used in [33]. As an alternative, which also satisfies the regularity needed, we propose the residual network also used for deterministic control in [27]. The network is given by

$$(3.1) \qquad \Phi(\boldsymbol{y}; \boldsymbol{\theta}) = \boldsymbol{w}^{\top}\mathcal{NN}(\boldsymbol{y}; \boldsymbol{\theta}_{\mathcal{NN}}) + \frac{1}{2}\boldsymbol{y}^{\top}(\boldsymbol{A}^{\top}\boldsymbol{A})\boldsymbol{y} + \boldsymbol{b}^{\top}\boldsymbol{y} + c,$$

with trainable weights $\boldsymbol{\theta} = (\boldsymbol{w}, \boldsymbol{\theta}_{\mathcal{NN}}, \boldsymbol{A}, \boldsymbol{b}, c)$. Here the inputs $\boldsymbol{y} = (s, \boldsymbol{z}(s)) \in \mathbb{R}^{d+1}$ correspond to time-space, $\mathcal{NN}(\boldsymbol{y}; \boldsymbol{\theta}_{\mathcal{NN}})\colon \mathbb{R}^{d+1} \to \mathbb{R}^m$ is a neural network, and $\boldsymbol{\theta}$ contains the trainable weights: $\boldsymbol{w} \in \mathbb{R}^m$, $\boldsymbol{\theta}_{\mathcal{NN}} \in \mathbb{R}^p$, $\boldsymbol{A} \in \mathbb{R}^{\gamma \times (d+1)}$, $\boldsymbol{b} \in \mathbb{R}^{d+1}$, $c \in \mathbb{R}$, where rank $\gamma = \min(10, d + 1)$ limits the number of parameters in $\boldsymbol{A}^{\top}\boldsymbol{A}$. Here, $\boldsymbol{A}$, $\boldsymbol{b}$, and c model quadratic potentials, i.e., linear dynamics; $\mathcal{NN}$ models nonlinear dynamics.

In our experiments, for $\mathcal{NN}$, we either use a MLP [11]

$$(3.2) \qquad \begin{aligned} \boldsymbol{a}_0 &= \mathrm{act}(\boldsymbol{K}_0\boldsymbol{y} + \boldsymbol{b}_0) \\ \boldsymbol{a}_{i+1} &= \mathrm{act}(\boldsymbol{K}_{i+1}\boldsymbol{a}_i + \boldsymbol{b}_{i+1}), \quad 0 \le i \le M - 2 \\ \mathcal{NN}(\boldsymbol{y}; \boldsymbol{\theta}_{\mathcal{NN}}) &= \mathrm{act}(\boldsymbol{K}_M\boldsymbol{a}_{M-1} + \boldsymbol{b}_M), \end{aligned}$$

or a residual neural network (ResNet) [14]

$$\boldsymbol{a}_0 = \mathrm{act}(\boldsymbol{K}_0 \boldsymbol{y} + \boldsymbol{b}_0)$$

(3.3) $$\boldsymbol{a}_{i+1} = \boldsymbol{a}_i + \mathrm{act}(\boldsymbol{K}_{i+1} \boldsymbol{a}_i + \boldsymbol{b}_{i+1}), \quad 0 \le \mathrm{i} \le \mathrm{M} - 2$$

$$\mathcal{NN}(\boldsymbol{y}; \boldsymbol{\theta}_{\mathcal{NN}}) = \boldsymbol{a}_{M-1} + \mathrm{act}(\boldsymbol{K}_{\mathrm{M}} \boldsymbol{a}_{M-1} + \boldsymbol{b}_{\mathrm{M}}),$$

with neural network weights $\boldsymbol{\theta}_{\mathcal{NN}} = (\boldsymbol{K}_0, \dots, \boldsymbol{K}_M, \boldsymbol{b}_0, \dots, \boldsymbol{b}_M)$ where $\boldsymbol{b}_i \in \mathbb{R}^m \ \forall i$, $\boldsymbol{K}_0 \in \mathbb{R}^{m \times (d+1)}$, and $\{\boldsymbol{K}_1, \dots, \boldsymbol{K}_M\} \in \mathbb{R}^{m \times m}$ with M being the depth of the network. The choice of the element-wise nonlinearity $\mathrm{act}(\cdot)$ is discussed in the respective experiments.

3.2. Training Problem. Ideally, we would choose $\boldsymbol{\theta}$ such that $\Phi(s, \boldsymbol{z}; \boldsymbol{\theta})$ is equal to the value function of the control problem globally, that is, for all $(s, \boldsymbol{z}) \in [t, T] \times \mathbb{R}^d$. Since this is known to be cursed by the dimensionality for reasonable problem sizes, we resort to a semi-global approach which enforces this property at randomly sampled points in the space time domain.

To generate samples, we first obtain initial states $\boldsymbol{x} \sim \rho$ from some (possibly Dirac) distribution ρ and then use an Euler Maruyama scheme with $N+1$ equidistant time points $s_0, \dots, s_N$ and step size $\mathrm{d}s = (T - t)/N$. This yields a state trajectory starting at $\boldsymbol{z}_0 = \boldsymbol{x}$ via

(3.4) $$\boldsymbol{z}_{i+1} = \boldsymbol{z}_i + f(s_i, \boldsymbol{z}_i, \boldsymbol{u}_i)\mathrm{d}s + \sigma(s_i, \boldsymbol{z}_i)\mathrm{d}\boldsymbol{W}_i, \quad i = 0, \dots, N-1$$

where $\mathrm{d}\boldsymbol{W}_i \sim \mathcal{N}(\boldsymbol{0}, \mathrm{d}s \cdot \boldsymbol{I}_d)$, and $\boldsymbol{u}_i = \boldsymbol{u}_{t,x}^*(s_i, \boldsymbol{z}_i)$ is the optimal control obtained from the feedback, that is, form (2.6)

$$\boldsymbol{u}_i^* \in \arg\max_{\boldsymbol{u} \in U} \mathcal{H}\left(s_i, \boldsymbol{z}_i, -\nabla\Phi(s_i, \boldsymbol{z}_i; \boldsymbol{\theta}), -\sigma(s_i, \boldsymbol{z}_i)^\top \nabla^2 \Phi(s_i, \boldsymbol{z}_i; \boldsymbol{\theta}), \boldsymbol{u}\right).$$

A few comments are in place. First, it is important to note that due to the feedback form, the sampled trajectories depend on the parameters of the value function. Second, the addition of this drift term, motivated by control theory, is the key difference to neural network solvers for the more general class of semi-linear elliptic PDEs [12, 33]. Third, the drift term can also be motivated by the fact that for $\sigma \to 0$, the trajectories above approximate the characteristic curves of the non-viscous HJB equation; thereby our SOC approach coincides with that for deterministic OC in [27].

To simplify our notation below, we collect the states, control, and noise along the discrete trajectories in (3.4) column-wise in the matrices

$$\boldsymbol{Z} \in \mathbb{R}^{d \times N}, \quad \boldsymbol{U} \in \mathbb{R}^{k \times N}, \quad \mathbf{dW} \in \mathbb{R}^{d \times N}.$$

To learn the parameters of the neural networks in an unsupervised way (that is, assuming neither analytic values of Φ nor optimal control trajectories), we approximately solve the minimization problem

(3.5) $$\min_{\boldsymbol{\theta}} \mathbb{E}_{\boldsymbol{x} \sim \rho} \left\{ \mathbb{E}_{\boldsymbol{Z}, \boldsymbol{U}, \mathbf{dW} | \boldsymbol{x}} \left\{ \beta_1 P_{\mathrm{BSDE}}(\boldsymbol{Z}, \boldsymbol{U}, \mathbf{dW}) + \beta_2 P_{\mathrm{HJB}}(\boldsymbol{Z}) + \beta_3 J(\boldsymbol{Z}, \boldsymbol{U}) \right. \right.$$

$$\left. \left. + \beta_4 |G(\boldsymbol{z}_N) - \Phi(s_N, \boldsymbol{z}_N; \boldsymbol{\theta})| + \beta_5 |\nabla G(\boldsymbol{z}_N) - \nabla\Phi(s_N, \boldsymbol{z}_N; \boldsymbol{\theta})| \right\} \right\},$$

where the terms in the objective function consist of penalty functions for violations of the BSDE system and the HJB PDE, the control objective, and penalty terms for the terminal condition, respectively, and are defined below. The relative influence of each term is controlled by the components of $\beta \in \mathbb{R}_+^5$. Different choices of β allow us to

experiment with different learning approaches; for example, setting $\beta_1 = \beta_4 = \beta_5 = 1$ and $\beta_2 = \beta_3 = 0$ provides the same loss function as in [33] while $\beta_1 = 0$ and $\beta_i > 0$, $i \in \{2, 3, 4, 5\}$ gives the loss function used for deterministic OC problems in [27].

We penalize the violation of the BSDE (2.13) via

$$(3.6)$$
$$P_{\mathrm{BSDE}}(\boldsymbol{Z}, \boldsymbol{U}, \mathbf{dW}) = \sum_{i=0}^{N-1} |\Phi_{i+1}(\boldsymbol{\theta}) - \Phi_i(\boldsymbol{\theta}) + L(s_i, \boldsymbol{z}_i, \boldsymbol{u}_i)\mathrm{d}s - \nabla\Phi_i(\boldsymbol{\theta})^\top \sigma(s_i, \boldsymbol{z}_i)\mathrm{d}\boldsymbol{W}_i|^p$$

where we use the abbreviations $\Phi_i(\boldsymbol{\theta}) := \Phi(s_i, \boldsymbol{z}_i; \boldsymbol{\theta})$ and $\nabla\Phi_i(\boldsymbol{\theta}) := \nabla\Phi(s_i, \boldsymbol{z}_i; \boldsymbol{\theta})$ and the exponent $p \in \{1, 2\}$. Similarly, the HJB penalty term reads

$$(3.7) \qquad P_{\mathrm{HJB}}(\boldsymbol{Z}) = \mathrm{d}s \sum_{i=1}^{N} |H(s_i, \boldsymbol{z}_i, -\nabla\Phi_i(\theta), -\sigma(s_i, \boldsymbol{z}_i)^\top \nabla^2\Phi_i(\theta)) - \partial_s\Phi_i(\boldsymbol{\theta})|^p,$$

where $\nabla^2\Phi_i(\theta) := \nabla^2\Phi(s_i, \boldsymbol{z}_i; \theta)$, $\partial_s\Phi_i(\theta) := \partial_s\Phi(s_i, \boldsymbol{z}_i; \theta)$. Finally, we approximate the objective functional via

$$J(\boldsymbol{Z}, \boldsymbol{U}) = G(\boldsymbol{z}_N) + \mathrm{d}s \sum_{i=1}^{N} L(s_i, \boldsymbol{z}_i, \boldsymbol{u}_i).$$

In principle, any stochastic approximation approach can be used to approximately solve the above optimization problem. Here, we use Adam and sample a minibatch of trajectories originating in n i.i.d. samples from ρ.

3.3. Implementation. We implement and test our proposed approach in two software environments.

To obtain a direct comparison with [33] we modify the FBSNN code accompanying the paper. To this end, we created a publicly available fork at https://github.com/EmoryMLIP/FBSNNs. Our main modifications are adding the control objective in the training loss and also adding the proposed drift to the forward dynamics. The experiments carried out here use a standard MLP as the neural network without the quadratic term added in (3.1).

In order to further simplify the experimentation, we also implement our own PyTorch code available at https://github.com/EmoryMLIP/NeuralHJB. Our implementation contains all loss terms in (3.5). We implement both sampling techniques: pure random walk and the proposed one informed by PMP. This facilitates comparisons of our approach with other available methods and simplifies developing new examples.

We tested most of our examples using either Intel Xeon E5-4627 CPU or Nvidia P100 GPU.

4. Numerical Experiments. We test the efficacy of our proposed algorithm on two different SOC problems. In subsection 4.1, we introduce a two-dimensional trajectory planning problem to visualize the difference between purely random exploration and our proposed sampling scheme. In subsection 4.2, we compare our approach to those in [7, 12] using a 100-dimensional benchmark problem. For the original version of this problem, our method shows faster initial convergence and time-to-solution but comparable accuracy. We modify the terminal cost of this problem to further highlight the importance of the feedback form in the sampling (see subsection 4.2.2) and experimentally show that the growth of network size is not exponential with increasing dimensions.

4.1. 2D Trajectory Planning Problem. To visualize the behaviour of our PMP-based sampling approach, we consider a two-dimensional test problem.

The problem consists of planning an optimal trajectory from the initial state $\boldsymbol{x} \sim \rho = \mathcal{N}((-1.5, -1.5)^\top, \boldsymbol{I}_2)$ to the target $\boldsymbol{x}_{\text{target}} = (1.5, 1.5)^\top$. To make the problem interesting, a hill is placed at the origin, which adds height-dependent cost for traveling around that region given by

$$Q(\boldsymbol{z}) = 50 \exp\left(-\frac{\|\boldsymbol{z}\|^2}{0.8}\right).$$

The dynamics for the problem read

(4.1) $$f(s, \boldsymbol{z}, \boldsymbol{u}) = \boldsymbol{u} \quad \text{and} \quad \sigma = 0.5.$$

The running cost and terminal cost of the problem are given respectively, by

(4.2) $$L(s, \boldsymbol{z}, \boldsymbol{u}) = \frac{1}{2}\|\boldsymbol{u}\|^2 + Q(\boldsymbol{z}) \quad \text{and} \quad G(\boldsymbol{z}) = 50 \cdot \|\boldsymbol{z} - \boldsymbol{x}_{\text{target}}\|^2.$$

The corresponding HJB equation is

(4.3) $$\partial_s \Phi(s, \boldsymbol{z}) + \frac{1}{8}\Delta\Phi(s, \boldsymbol{z}) - \frac{1}{2}\|\nabla\Phi(s, \boldsymbol{z})\|^2 + Q(\boldsymbol{z}) = 0.$$

With the above calculations, the sampling SDE (3.4) simplifies to

(4.4) $$\boldsymbol{z}_{i+1} = \boldsymbol{z}_i - \nabla\Phi(s_i, \boldsymbol{z}_i)\mathrm{d}s + \frac{1}{2}\mathrm{d}\boldsymbol{W}_i.$$

We approximate the value function using the model in (3.1) where $\mathcal{NN}$ is a two-layer residual neural network with 32 neurons per layer and the rank of the quadratic term is $\gamma = 3$. Overall, the model has 1229 trainable weights. As activation function we chose the anti-derivative of the tanh function as in [27]. We use the penalty parameters $\beta = (1, 1, 1, 0.1, 0.1)$, that is, we enable both the penalty terms, P_{BSDE} and P_{HJB}. To approximately solve (3.5) we use 3000 steps of the Adam optimizer and batch size of 64. After 1500 iterations, we reduce the learning rate from 0.01 to 0.001. The average cost per iteration is about 0.1s using the CPU.

In Figure 1, we visualize the estimate of the value function and the optimal control policy at $t = 0$ as well as example trajectories for randomly chosen initial states after training the model. As expected, the trajectories travel from the initial points to the target while avoiding the obstacle in the center of the domain.

To gain more insight into the sampling, we store all states visited during training and plot them as two-dimensional histograms for different time points (left to right) in Figure 2. We compare the proposed PMP-based sampling (Figure 2b) to the pure random walk (i.e., without drift) used in [12, 33] (Figure 2a). As to be expected, sampling without drift leads to visiting points around the initial states in all (even suboptimal) directions. It can also be seen that there are almost no samples close to the target. With the drift term and after training, the sampled states cover more likely paths between the initial and target states.

Another way to interpret the histogram plots in Figure 2 is that they show the semi-global nature of neural network approaches for SOC problems and, crucially, that the regions of the state space visited in training can vary a lot. Since the loss function in (3.5) penalizes the HJB and BSDE loss only around points sampled from the forward SDE, one would expect that the trained model is more reliable in regions that are frequently visited.
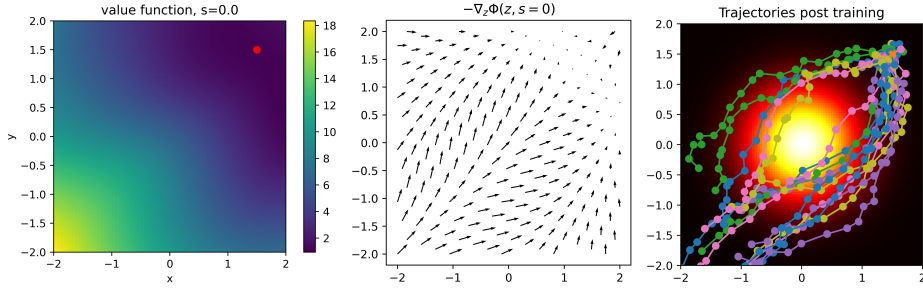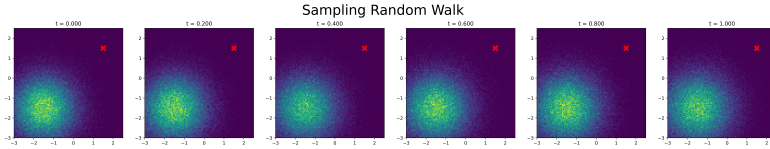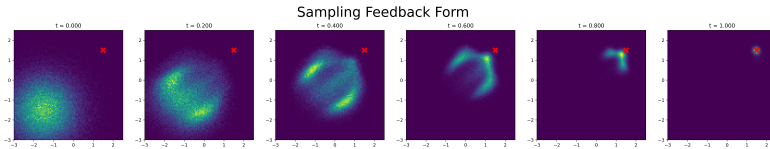
Fig. 1: Results of the two-dimensional test problem. Left: Value function function estimate $\Phi(0, \cdot; \boldsymbol{\theta}^*)$ after training. Middle: Quiver plot of optimal controls at $s = 0$. Right: Trajectories from randomly chosen initial states.

(a) Training samples with pure random walk as FSDE as also used in [12] and [33].

(b) Training samples with PMP-based drift term.

Fig. 2: We visualize training samples of a pure random walk sampler (top row) and our proposed PMP-based sampler (bottom row) for the two-dimensional test problem. At six time points (left to right), we visualize the sampled states as two-dimensional histrograms. As expected, the pure random walk explores the area around the initial state in all (even suboptimal) directions, while the proposed approach learns to sample around approximately optimal trajectories.

4.2. 100-dimensional example. We consider the 100-dimensional benchmark SOC problem also used in [7, 12] with initial state $\boldsymbol{x} = (0, 0, \ldots, 0)^\top \in \mathbb{R}^{100}$ corresponding to time $t = 0$. The drift and diffusion of the system are given by

$$f(s, \boldsymbol{z}, \boldsymbol{u}) = 2\boldsymbol{u} \quad \text{and} \quad \sigma = \sqrt{2},$$

respectively. The terminal and Lagrangian cost are

$$(4.5) \qquad G(\boldsymbol{z}) = \ln\left(\frac{1 + \|\boldsymbol{z}\|^2}{2}\right), \quad \text{and} \quad L(s, \boldsymbol{z}, \boldsymbol{u}) = \|\boldsymbol{u}\|^2,$$

respectively. We compute the Hamiltonian (2.4) as

$$H(s, \boldsymbol{z}, \boldsymbol{p}, \boldsymbol{M}) = \sup_{\boldsymbol{u} \in U} \left\{ \frac{\sigma}{2} \mathrm{tr}\,(\boldsymbol{M}) + \boldsymbol{p} \cdot f(s, \boldsymbol{z}, \boldsymbol{u}) - L(s, \boldsymbol{z}, \boldsymbol{u}) \right\}$$
$$= \sup_{\boldsymbol{u} \in U} \left\{ \frac{1}{\sqrt{2}} \mathrm{tr}\,(\boldsymbol{M}) + \boldsymbol{p} \cdot 2\boldsymbol{u} - \|\boldsymbol{u}\|^2 \right\}$$

Using the first-order necessary condition we get

$$0 = 2\boldsymbol{u} - 2\boldsymbol{p} \implies \boldsymbol{u} = \boldsymbol{p},$$

and using this closed form for $\boldsymbol{u}$, the Hamiltonian is given by

$$H(s, \boldsymbol{z}, \boldsymbol{p}, \boldsymbol{M}) = \frac{1}{\sqrt{2}} \mathrm{tr}\,(\boldsymbol{M}) + \|\boldsymbol{p}\|^2.$$

Hence, the HJB equation satisfied by the value function, $\Phi(\cdot, \cdot)$, reads

$$(4.6) \qquad \frac{\partial}{\partial s} \Phi(s, \boldsymbol{z}) + \Delta \Phi(s, \boldsymbol{z}) - \|\nabla \Phi(s, \boldsymbol{z})\|^2 = 0, \quad \Phi(T, \boldsymbol{z}) = G(\boldsymbol{z}).$$

and its solution is given by

$$(4.7) \qquad \Phi(s, \boldsymbol{z}) = -\ln\left(\mathbb{E}\left(\exp\left(-G\left(\boldsymbol{z} + \sqrt{2}\, \mathrm{d}W(T - s) \right) \right) \right) \right),$$

which we use to test the performance of our method.

Finally, we note that the forward SDE (3.4) we propose to use for sampling the state space simplifies to

$$(4.8) \qquad \boldsymbol{z}_{i+1} = \boldsymbol{z}_i - 2 \nabla_{\boldsymbol{z}} \Phi(s_i, \boldsymbol{z}_i)\mathrm{d}s + \sqrt{2}\mathrm{d}\boldsymbol{W}_i.$$

4.2.1. The importance of sampling. To demonstrate the impact of using the feedback form to sample the state space, we use the same neural network model as in [33], which is given by a 5-layer feed forward neural network with 256 neurons per hidden layer to approximate the solution $\Phi(s, \boldsymbol{z})$. We partition the time interval $[0, 1]$ using 50-uniformly spaced points. We use the same penalty parameters as in the original code, that is, $\beta = (1, 0, 20, 1, 1)$. We use the Adam optimizer [18] to update the parameters of the network with a batch size of 64 using 50k iterations. The average cost per 100 iterations was 27s using the CPU. For the following experiments, we use (3.5) excluding P_{HJB} penalty and compare our method with FBSNNs in [33].

Method	20k iterations		50k iterations	
	RE	RE$_0$	RE	RE$_0$
FBSNN	0.54%	0.12%	0.39%	0.045%
Ours	0.48%	0.0083%	0.39%	0.012%

Table 1: Relative errors for (4.6) obtained using our method and method in [33]

In Figure 3 we plot the exact solution (black-dashed line) (4.7), the learned solution using our approach (blue-solid line) and the solution learned using FBSNNs
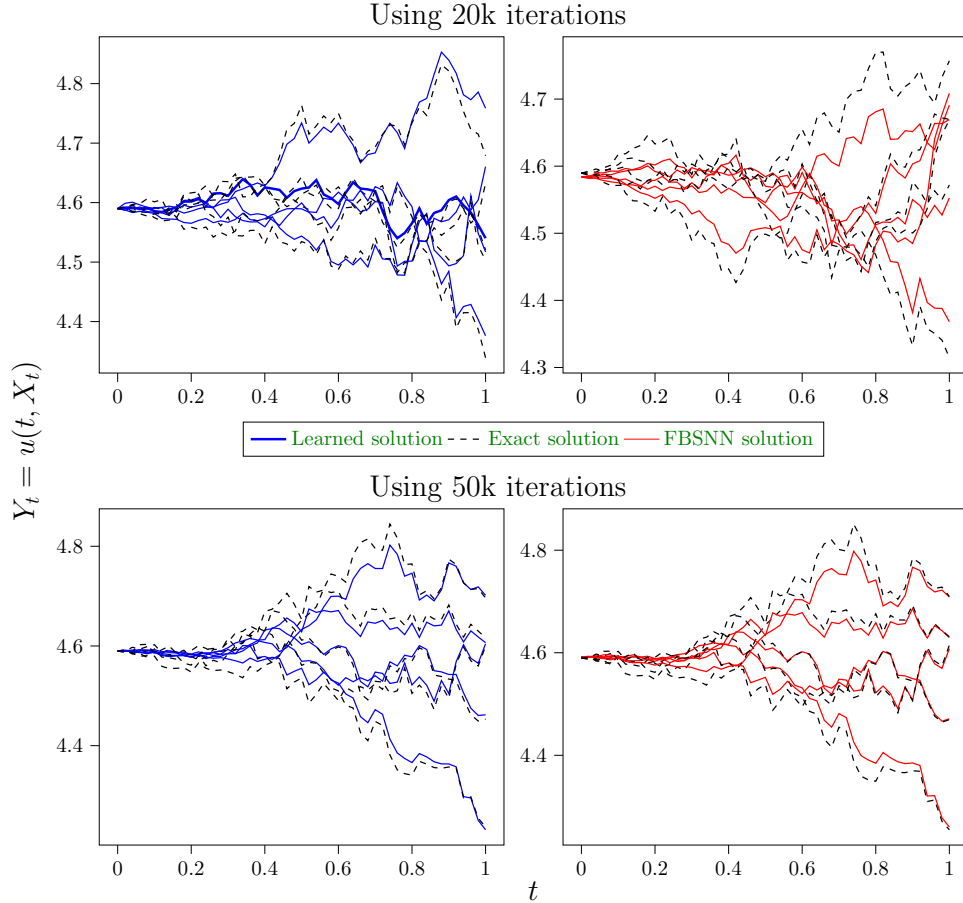
Fig. 3: Solution to (4.6) obtained using our method and method in [33]

(red-solid line) along five random trajectories. In the top row, we present the results obtained after training the networks for 20K iterations with a learning rate of 10^{-3} and the bottom row presents the results after training the networks for 20K and 30K iterations with learning rates 10^{-3} and 10^{-4}, respectively. These results suggest that our approach is approximating the value function better, especially in early iterations, as compared to the FBSNNs.

In Table 1, we also compare the learned solutions to the exact solution Φ in (4.7) by computing the average relative errors,

$$RE = \frac{\|\Phi(\cdot, \cdot; \boldsymbol{\theta}) - \Phi(\cdot, \cdot)\|_2}{\|\Phi\|_2}, \quad RE_0 = \frac{|\Phi(0, \boldsymbol{z}(0); \boldsymbol{\theta}) - \Phi(0, \boldsymbol{z}(0))|}{|\Phi(0, \boldsymbol{z}(0))|},$$

for ten random trajectories. Our method attains lower errors, especially for the initial values and at the earlier iterations.

4.2.2. Shifted target. In the example above, the minimizer of the terminal function coincides with the initial state $\boldsymbol{x} = (0, 0, \ldots, 0)^{\top}$. Therefore, even a random walk without drift (as used in [33, 12]) will sample around the optimal terminal state,
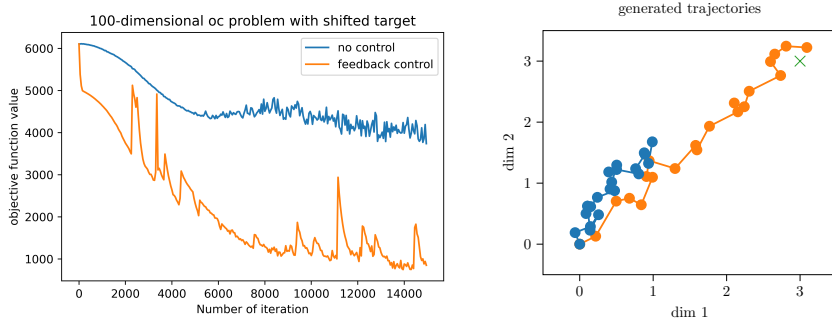
Fig. 4: Computational results for the modified 100-dimensional benchmark problem in subsection 4.2.2. Left: Control objective for both methods given the same initial state, blue line represents results using FBSNNs in [33], and orange line denotes our method. Right: Trajectory examples generated using learned value functions on two randomly selected dimensions. Orange line represents our method and blue line FBSNNs.

which is critical to accurately approximate the value function. This also means that after training using our approach, the drift term in the sampler is relatively small and that the above experiment does not fully show the advantages of our method.

To shed more light on the importance of sampling, we modify the terminal cost to

$$G(z) = 1000 \ln \left(\frac{1 + \|z - z_{\text{target}}\|^2}{2} \right),$$

with $z_{\text{target}} = [3, 3, \ldots, 3]^T$, so that the target for the state variable z at final time T no longer coincides with the initial state. Similar to the two-dimensional test problem in subsection 4.1, solving the modified problem now requires sampling around the target and we expect to benefit from the added drift term.

We compare our method to FBSNNs on the modified problem keeping the same network structure and hyper-parameters. We use a smaller $\sigma = \frac{2\sqrt{2}}{5}$ to improve training speed. We evaluate the performance of the methods using the objective functional J defined in (2.2) at the control obtained from the feedback form via the respective value function approximations. For this experiment we use GPU to train and the results of this comparison are shown in Figure 4.

To reduce the effect of the Brownian motion, we run the experiments for each method on the same problem five times and plot the average values corresponding to training iterations, we also further reduced the effect of the random walk in evaluation for clearer visualization.

As can be seen in Figure 4 (left), our method not only yields faster initial convergence but also achieves a considerably lower control objective. This indicates that the controls obtained from our approach are more effective, that is, they are closer to optimal. Since FBSNNs use a Brownian motion with no drift, the sampling is unlikely to discover the target. Consequently, the generated trajectories in Figure 4 (right) from our method approximately reach the target, while the trajectories obtained from the FBSNN method stay closer to the initial state.
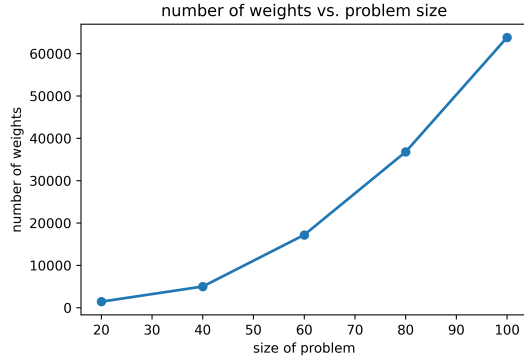
Fig. 5: We compare the necessary number of weights required to train a working model for problem with different sizes. Each problem follows the same setup used in subsection 4.2.2 with the only difference being the dimensionality of the problem. For a model to be considered working we require that under the feedback form the expected value for the final state of the agent to be within 0.02 range of the target.

4.2.3. Scalability with problem dimension. As some evidence that the proposed NN approach can mitigate the curse of dimensionality, we vary the dimension in $\{20, 40, 60, 80, 100\}$ in subsection 4.2.2. We use the same network architecture from the previous session. For each dimension, we increase the width of the network until the training yields a network whose trajectories end sufficiently close to the target.

With our method we notice that the number of weights required to solve the problem does not increase exponentially with dimensionality, suggesting that our method does ease the CoD. It is however worth mentioning that this experiment does not aim to determine the minimum number of weights required to solve each problem.

5. Discussion. We propose a neural network approach for approximately solving Hamilton-Jacobi-Bellman PDEs arising in high-dimensional stochastic optimal control. Similar to existing approaches [33, 12], we parameterize the value function with a neural network and we experiment with different losses to train the weights of the network. What sets our work apart from existing approaches is the use of feedback form given by the stochastic Pontryagin maximum principle to design the forward SDE used to explore the state space during training.

Using an intuitive two-dimensional test problem, we visualize that the improved sampling strategy allows us to effectively learn the value function and determine the relevant regions of the state space; see subsection 4.1. Based on this insight, we modify the 100-dimensional test problem also used in [33, 12] by shifting the minimizer of the terminal costs; see subsection 4.2.2. Thereby, we demonstrate that our proposed method dramatically improves the quality of the obtained control.

In additional numerical experiments, we show that the modified forward SDE and the control objective can lead to faster initial convergence compared to the approaches in [33, 12] (see subsection 4.2) and show that, for the problem at hand, the network size needed to achieve a given tolerance does not need to grow exponentially.

We refer to our method as a semi-global method for solving the HJB equation since we do not aim at estimating the value function accurately globally. Instead, we seek to approximate the value function well for states likely to be visited by

optimal trajectories of the SOC problem. In theory, any point in the state space has a positive probability to be visited due to the stochasticity in the dynamics. However, the histogram plots in Figure 2 suggest that the density of the optimal trajectory is concentrated in a small subset of the state space. Therefore, focussing the exploration to this subsets may have practical advantages over FSDEs that are pure random walks.

Another benefit of our proposed forward SDE compared to purely random exploration is that it coincides with the characteristic curves of the HJB equation as the stochasticity of the system is reduced. Therefore, our work can be seen as an extension of the neural network approaches for deterministic control problems in [27].

Compared to neural network approaches for semi-linear elliptic/parabolic PDEs such as [33, 12] it is important to highlight that our approach is only applicable to HJB equations arising in stochastic optimal control. Since our forward SDE is derived from optimality principles, extending it to other high-dimensional PDEs (e.g., Black Scholes and Allen Cahan equations) is not obvious and may be impossible.

In future work, we will apply our approach to problems with non-constant and non-scalar diffusion coefficients. Although the theoretical framework supports this, we are not aware of any practical algorithms for this case. To overcome the challenge of efficiently computing Hessians of the value function needed for HJB penalties, we will use the hessQuik package [26].

REFERENCES

[1] F. ANTONELLI, *Backward-forward stochastic differential equations*, Ann. Appl. Probab., 3 (1993), pp. 777–793, http://links.jstor.org/sici?sici=1050-5164(199308)3:3⟨777:BSDE⟩2.0. CO;2-5&origin=MSN.

[2] R. BELLMAN, *Dynamic programming and stochastic control processes*, Information and Control, 1 (1958), pp. 228–239.

[3] C. BENDER AND R. DENK, *A forward scheme for backward SDEs*, Stochastic Process. Appl., 117 (2007), pp. 1793–1812, https://doi.org/10.1016/j.spa.2007.03.005, https://doi.org/10. 1016/j.spa.2007.03.005.

[4] D. BERTSEKAS AND J. TSITSIKLIS, *Neuro-dynamic programming: an overview*, in Proceedings of 1995 34th IEEE Conference on Decision and Control, vol. 1, 1995, pp. 560–564 vol.1, https://doi.org/10.1109/CDC.1995.478953.

[5] P. CHERIDITO, H. M. SONER, N. TOUZI, AND N. VICTOIR, *Second-order backward stochastic differential equations and fully nonlinear parabolic PDEs*, Comm. Pure Appl. Math., 60 (2007), pp. 1081–1110, https://doi.org/10.1002/cpa.20168, https://doi.org/10.1002/cpa. 20168.

[6] H. DONG AND N. V. KRYLOV, *The rate of convergence of finite-difference approximations for parabolic Bellman equations with Lipschitz coefficients in cylindrical domains*, Appl. Math. Optim., 56 (2007), pp. 37–66, https://doi.org/10.1007/s00245-007-0879-4, https: //doi.org/10.1007/s00245-007-0879-4.

[7] W. E, J. HAN, AND A. JENTZEN, *Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations*, Commun. Math. Stat., 5 (2017), pp. 349–380, https://doi.org/10.1007/s40304-017-0117-6, https://doi.org/10.1007/s40304-017-0117-6.

[8] I. EXARCHOS, E. THEODOROU, AND P. TSIOTRAS, *Stochastic differential games: a sampling approach via FBSDEs*, Dyn. Games Appl., 9 (2019), pp. 486–505, https://doi.org/10.1007/ s13235-018-0268-4, https://doi.org/10.1007/s13235-018-0268-4.

[9] I. EXARCHOS AND E. A. THEODOROU, *Stochastic optimal control via forward and backward stochastic differential equations and importance sampling*, Automatica J. IFAC, 87 (2018), pp. 159–165, https://doi.org/10.1016/j.automatica.2017.09.004, https://doi.org/10.1016/j. automatica.2017.09.004.

[10] W. H. FLEMING AND H. M. SONER, *Controlled Markov Processes and Viscosity Solutions*, vol. 25 of Stochastic Modelling and Applied Probability, Springer, New York, second ed., 2006.

[11] I. GOODFELLOW, Y. BENGIO, AND A. COURVILLE, *Deep Learning*, MIT Press, 2016. http: //www.deeplearningbook.org.

[12] J. HAN, A. JENTZEN, AND W. E, *Solving high-dimensional partial differential equations using deep learning*, Proc. Natl. Acad. Sci. USA, 115 (2018), pp. 8505–8510, https://doi.org/10.1073/pnas.1718942115, https://doi.org/10.1073/pnas.1718942115.

[13] K. P. HAWKINS, A. PAKNIYAT, E. THEODOROU, AND P. TSIOTRAS, *Forward-backward rapidly-exploring random trees for stochastic optimal control*, 2020, https://doi.org/10.48550/ARXIV.2006.12444, https://arxiv.org/abs/2006.12444.

[14] K. HE, X. ZHANG, S. REN, AND J. SUN, *Deep residual learning for image recognition*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2016.

[15] C. HURÉ, H. PHAM, AND X. WARIN, *Deep backward schemes for high-dimensional nonlinear PDEs*, Math. Comp., 89 (2020), pp. 1547–1579, https://doi.org/10.1090/mcom/3514, https://doi.org/10.1090/mcom/3514.

[16] E. R. JAKOBSEN, *On the rate of convergence of approximation schemes for Bellman equations associated with optimal stopping time problems*, Math. Models Methods Appl. Sci., 13 (2003), pp. 613–644, https://doi.org/10.1142/S0218202503002660, https://doi.org/10.1142/S0218202503002660.

[17] S. JI, S. PENG, Y. PENG, AND X. ZHANG, *Three algorithms for solving high-dimensional fully coupled fbsdes through deep learning*, IEEE Intelligent Systems, 35 (2020), pp. 71–84, https://doi.org/10.1109/MIS.2020.2971597.

[18] D. P. KINGMA AND J. BA, *Adam: A method for stochastic optimization*, arXiv preprint arXiv:1412.6980, (2014).

[19] P. E. KLOEDEN AND E. PLATEN, *Numerical solution of stochastic differential equations*, vol. 23 of Applications of Mathematics (New York), Springer-Verlag, Berlin, 1992, https://doi.org/10.1007/978-3-662-12616-5, https://doi.org/10.1007/978-3-662-12616-5.

[20] N. V. KRYLOV, *The rate of convergence of finite-difference approximations for Bellman equations with Lipschitz coefficients*, Appl. Math. Optim., 52 (2005), pp. 365–399, https://doi.org/10.1007/s00245-005-0832-3, https://doi.org/10.1007/s00245-005-0832-3.

[21] K. KUNISCH AND D. WALTER, *Semiglobal optimal feedback stabilization of autonomous systems via deep neural network approximation*, ESAIM Control Optim. Calc. Var., 27 (2021), pp. Paper No. 16, 59, https://doi.org/10.1051/cocv/2021009, https://doi.org/10.1051/cocv/2021009.

[22] H. J. KUSHNER, *Numerical methods for stochastic control problems in continuous time*, SIAM J. Control Optim., 28 (1990), pp. 999–1048, https://doi.org/10.1137/0328056, https://doi.org/10.1137/0328056.

[23] H. J. KUSHNER AND P. DUPUIS, *Numerical methods for stochastic control problems in continuous time*, vol. 24 of Applications of Mathematics (New York), Springer-Verlag, New York, second ed., 2001, https://doi.org/10.1007/978-1-4613-0007-6, https://doi.org/10.1007/978-1-4613-0007-6. Stochastic Modelling and Applied Probability.

[24] J. MA, P. PROTTER, AND J. M. YONG, *Solving forward-backward stochastic differential equations explicitly—a four step scheme*, Probab. Theory Related Fields, 98 (1994), pp. 339–359, https://doi.org/10.1007/BF01192258, https://doi.org/10.1007/BF01192258.

[25] J. MA AND J. YONG, *Forward-backward stochastic differential equations and their applications*, vol. 1702 of Lecture Notes in Mathematics, Springer-Verlag, Berlin, 1999.

[26] E. NEWMAN AND L. RUTHOTTO, *'hessquik': Fast hessian computation of composite functions*, Journal of Open Source Software, 7 (2022), p. 4171, https://doi.org/10.21105/joss.04171, https://doi.org/10.21105/joss.04171.

[27] D. ONKEN, L. NURBEKYAN, X. LI, S. W. FUNG, S. OSHER, AND L. RUTHOTTO, *A neural network approach for high-dimensional optimal control*, 2021, https://arxiv.org/abs/2104.03270.

[28] E. PARDOUX AND S. TANG, *Forward-backward stochastic differential equations and quasilinear parabolic PDEs*, Probab. Theory Related Fields, 114 (1999), pp. 123–150, https://doi.org/10.1007/s004409970001, https://doi.org/10.1007/s004409970001.

[29] M. PEREIRA, Z. WANG, T. CHEN, E. REED, AND E. THEODOROU, *Feynman-kac neural network architectures for stochastic control using second-order fbsde theory*, in Proceedings of the 2nd Conference on Learning for Dynamics and Control, A. M. Bayen, A. Jadbabaie, G. Pappas, P. A. Parrilo, B. Recht, C. Tomlin, and M. Zeilinger, eds., vol. 120 of Proceedings of Machine Learning Research, PMLR, 10–11 Jun 2020, pp. 728–738, https://proceedings.mlr.press/v120/pereira20a.html.

[30] H. PHAM, *Continuous-time stochastic control and optimization with financial applications*, vol. 61 of Stochastic Modelling and Applied Probability, Springer-Verlag, Berlin, 2009, https://doi.org/10.1007/978-3-540-89500-8, https://doi.org/10.1007/978-3-540-89500-8.

[31] L. S. PONTRYAGIN, V. G. BOLTYANSKII, R. V. GAMKRELIDZE, AND E. F. MISHCHENKO, *The Mathematical Theory of Optimal Processes*, Translated by K. N. Trirogoff; edited by L.

W. Neustadt, Interscience Publishers John Wiley & Sons, Inc. New York-London, 1962.

[32] W. B. Powell, *Approximate dynamic programming*, Wiley Series in Probability and Statistics, Wiley-Interscience [John Wiley & Sons], Hoboken, NJ, 2007, https://doi.org/10.1002/9780470182963, https://doi.org/10.1002/9780470182963. Solving the curses of dimensionality.

[33] M. Raissi, *Forward-backward stochastic neural networks: Deep learning of high-dimensional partial differential equations*, arXiv preprint arXiv:1804.07010, (2018).

[34] R. F. Stengel, *Optimal control and estimation*, Dover Publications, Inc., New York, 1994. Corrected reprint of the 1986 original.

[35] J. Wang and P. A. Forsyth, *Maximal use of central differencing for Hamilton-Jacobi-Bellman PDEs in finance*, SIAM J. Numer. Anal., 46 (2008), pp. 1580–1601, https://doi.org/10.1137/060675186, https://doi.org/10.1137/060675186.

[36] S. Yensiri and R. J. Skulkhu, *An investigation of radial basis function-finite difference (rbf-fd) method for numerical solution of elliptic partial differential equations*, Mathematics, 5 (2017), https://doi.org/10.3390/math5040054, https://www.mdpi.com/2227-7390/5/4/54.

[37] J. Yong and X. Y. Zhou, *Stochastic controls*, vol. 43 of Applications of Mathematics (New York), Springer-Verlag, New York, 1999, https://doi.org/10.1007/978-1-4612-1466-3, https://doi.org/10.1007/978-1-4612-1466-3. Hamiltonian systems and HJB equations.