Deep splitting method for parabolic PDEs

Beck, Christian; Becker, Sebastian; Cheridito, Patrick; Jentzen, Arnulf; Neufeld, Ariel

2021

Beck, C., Becker, S., Cheridito, P., Jentzen, A. & Neufeld, A. (2021). Deep splitting method for parabolic PDEs. SIAM Journal On Scientific Computing, 43(5), A3135-A3154.

<https://dx.doi.org/10.1137/19M1297919>

<https://hdl.handle.net/10356/153744>

<https://doi.org/10.1137/19M1297919>

© 2021 Society for Industrial and Applied Mathematics. All rights reserved. This paper was published in SIAM Journal on Scientific Computing and is made available with permission of Society for Industrial and Applied Mathematics.

Downloaded on 26 Apr 2023 21:57:46 SGT

DEEP SPLITTING METHOD FOR PARABOLIC PDEs*

CHRISTIAN BECK[†], SEBASTIAN BECKER[†], PATRICK CHERIDITO[†],
ARNULF JENTZEN[‡], AND ARIEL NEUFELD[§]

Abstract. In this paper, we introduce a numerical method for nonlinear parabolic partial differential equations (PDEs) that combines operator splitting with deep learning. It divides the PDE approximation problem into a sequence of separate learning problems. Since the computational graph for each of the subproblems is comparatively small, the approach can handle extremely high dimensional PDEs. We test the method on different examples from physics, stochastic control, and mathematical finance. In all cases, it yields very good results in up to 10,000 dimensions with short run times.

Key words. nonlinear partial differential equations, splitting-up method, neural networks, deep learning

AMS subject classifications. 35K15, 65C05, 65M22, 65M75, 91G20, 93E20

DOI. 10.1137/19M1297919

1. Introduction. In this paper, we derive a numerical scheme for parabolic partial differential equations (PDEs) of the form

$$(1.1) \quad \frac{\partial}{\partial t} u(t, x) = F(x, u(t, x), \nabla_x u(t, x)) + \frac{1}{2} \text{Trace}(\sigma(x)\sigma^*(x) \text{Hess}_x u(t, x)),$$

$(t, x) \in (0, T] \times \mathbb{R}^d$, with initial condition $u(0, x) = \varphi(x)$, where $F: \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d \rightarrow \mathbb{R}$, $\sigma: \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$, and $\varphi: \mathbb{R}^d \rightarrow \mathbb{R}$ are appropriate continuous functions. Such PDEs describe various phenomena in nature, engineering, economics, and finance. They typically do not admit closed form solutions and, therefore, have to be solved numerically. In some applications, the dimension d can be high. For instance, in physics and engineering applications, $x \in \mathbb{R}^d$ typically models the coordinates of all components of a given system, whereas in derivative pricing and optimal investment problems, d usually corresponds to the number of underlying assets. Many classical PDEs, such as the standard heat and Black–Scholes equations, are linear. Using the Feynman–Kac representation, their solutions can efficiently be approximated in high dimensions with simple Monte Carlo averages. But if constraints or frictions are taken into account, or the PDE describes a control problem, the function F is no longer linear and (1.1) becomes much more challenging to solve for large d .

Numerical methods for PDEs have a long history. Classical approaches like finite differences and finite elements (see, e.g., [13, 69, 93]) are deterministic. In their

*Submitted to the journal’s Methods and Algorithms for Scientific Computing section November 5, 2019; accepted for publication (in revised form) July 6, 2021; published electronically September 13, 2021.

<https://doi.org/10.1137/19M1297919>

Funding: This work was supported by Swiss National Science Foundation grant 200020_175699 “Higher order numerical approximation methods for stochastic partial differential equations,” by the Deutsche Forschungsgemeinschaft under Germany’s Excellence Strategy EXC 2044-390685587, Mathematics Münster: Dynamics - Geometry - Structure, and by Nanyang Assistant Professorship grant “Machine Learning based Algorithms in Finance and Insurance.”

[†]Department of Mathematics, ETH Zurich, Zurich, 8092, Switzerland (christian.beck@math.ethz.ch, sebastian.becker@math.ethz.ch, patrick.cheridito@math.ethz.ch).

[‡]Institute for Analysis and Numerics, Faculty of Mathematics and Computer Science, University of Münster, Münster, 48149, Germany (ajentzen@uni-muenster.de).

[§]Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore 637371 (ariel.neufeld@ntu.edu.sg).

standard form, they work well for $d = 1, 2$, and 3 , but their complexity grows exponentially in d . To tackle higher dimensional problems, different simulation-based approaches have been developed that exploit a stochastic representation of the solution of the PDE. For instance, [1, 7, 8, 12, 15, 16, 17, 18, 19, 20, 21, 22, 24, 26, 35, 36, 37, 38, 39, 40, 53, 68, 71, 72, 75, 76, 77, 82, 83, 85, 87, 88, 89, 94, 96] use backward stochastic differential equation (BSDE) representations of PDEs and study approximation methods based on recursive polynomial regressions, [49, 51, 52, 80, 91, 95] investigate methods based on branching diffusion processes, and [28, 29, 56, 57, 58] analyze full-history recursive multilevel Picard methods. Recently, numerical methods for high-dimensional PDEs based on the idea to reformulate the PDE as a stochastic learning problem have been proposed in [27, 47]. This opens the door to the application of deep learning; see, e.g., [2, 3, 6, 10, 14, 30, 32, 33, 41, 48, 50, 54, 60, 73, 74, 78, 86, 90] for modifications and extensions. There are also already a few papers studying the convergence of deep learning based approximation methods for PDEs. For instance, [48, 90] derive convergence results without information on the convergence speed, whereas [11, 31, 43, 55, 61, 67] provide convergence and tractability results with dimension-independent convergence rates and error constants depending polynomially on the dimension.

In this paper, we develop a new deep learning method for parabolic PDEs that splits the differential operator into a linear and a nonlinear part. More precisely, we write

$$(1.2) \quad F(x, u(t, x), \nabla_x u(t, x)) = \langle \mu(x), \nabla_x u(t, x) \rangle_{\mathbb{R}^d} + f(x, u(t, x), \nabla_x u(t, x))$$

for suitable continuous functions $\mu: \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $f: \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d \rightarrow \mathbb{R}$. This decomposition is not unique. But the idea is that μ is chosen such that the nonlinearity $f(x, u(t, x), \nabla_x u(t, x))$ becomes small. Then we solve the PDE iteratively over small time intervals by approximating $f(x, u(t, x), \nabla_x u(t, x))$ and using the Feynman–Kac representation locally. This requires a recursive computation of conditional expectations. We approximate them by formulating them as minimization problems that can be approached with deep learning. This decomposes the PDE approximation problem into a sequence of separate learning problems. Since the computational graph for each of the subproblems is comparatively small, the method works for very high dimensional problems.

The rest of the paper is organized as follows. In section 2, we introduce the framework and derive the deep splitting method. In section 3, we test the approach on five different classes of high-dimensional nonlinear PDEs: Hamilton–Jacobi–Bellman (HJB) equations, nonlinear Black–Scholes equations, Allen–Cahn-type equations, nonlinear heat equations, and sine-Gordon-type equations.

2. Derivation of the proposed approximation algorithm. Fix $T \in (0, \infty)$ and $d \in \mathbb{N}$. Consider two at most polynomially growing continuous functions $\varphi: \mathbb{R}^d \rightarrow \mathbb{R}$ and $f: \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d \rightarrow \mathbb{R}$ together with two Lipschitz continuous functions $\mu: \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\sigma: \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$. Assume $u: [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ is an at most polynomially growing continuous function that is $C^{1,2}$ on $(0, T] \times \mathbb{R}^d$ and satisfies the PDE

$$(2.1) \quad \begin{aligned} \frac{\partial}{\partial t} u(t, x) &= f(x, u(t, x), \nabla_x u(t, x)) + \langle \mu(x), \nabla_x u(t, x) \rangle_{\mathbb{R}^d} \\ &\quad + \frac{1}{2} \text{Trace}(\sigma(x) \sigma^*(x) \text{Hess}_x u(t, x)), \end{aligned}$$

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

2.1. Temporal discretization. To approximate the solution u of the PDE (2.1), we discretize the equation in time and use a splitting-up method (see, e.g., [42, 44, 45]) to obtain a semidiscrete approximation problem. To do this, we choose $N \in \mathbb{N}$ and let $t_0, t_1, \dots, t_N \in [0, T]$ be real numbers such that

$$(2.2) \quad 0 = t_0 < t_1 < \dots < t_N = T.$$

Under appropriate integrability assumptions, it follows from (2.1) that for every $t \in [0, T]$, $x \in \mathbb{R}^d$ we have

$$(2.3) \quad \begin{aligned} u(t, x) &= \varphi(x) + \int_0^t f(x, u(s, x), \nabla_x u(s, x)) \, ds \\ &\quad + \int_0^t \left[\frac{1}{2} \text{Trace}(\sigma(x)\sigma^*(x) \text{Hess}_x u(s, x)) + \langle \mu(x), \nabla_x u(s, x) \rangle_{\mathbb{R}^d} \right] ds. \end{aligned}$$

In particular, for all $n \in \{0, 1, \dots, N-1\}$, $t \in [t_n, t_{n+1}]$, and $x \in \mathbb{R}^d$,

$$(2.4) \quad \begin{aligned} u(t, x) &= u(t_n, x) + \int_{t_n}^t f(x, u(s, x), (\nabla_x u)(s, x)) \, ds \\ &\quad + \int_{t_n}^t \left[\frac{1}{2} \text{Trace}(\sigma(x)\sigma^*(x) \text{Hess}_x u(s, x)) + \langle \mu(x), \nabla_x u(s, x) \rangle_{\mathbb{R}^d} \right] ds, \end{aligned}$$

and therefore,

$$(2.5) \quad \begin{aligned} u(t, x) &\approx u(t_n, x) + \int_{t_n}^{t_{n+1}} f(x, u(t_n, x), \nabla_x u(t_n, x)) \, ds \\ &\quad + \int_{t_n}^t \left[\frac{1}{2} \text{Trace}(\sigma(x)\sigma^*(x) \text{Hess}_x u(s, x)) + \langle \mu(x), \nabla_x u(s, x) \rangle_{\mathbb{R}^d} \right] ds, \end{aligned}$$

which can be written as

$$(2.6) \quad \begin{aligned} u(t, x) &\approx u(t_n, x) + f(x, u(t_n, x), \nabla_x u(t_n, x)) (t_{n+1} - t_n) \\ &\quad + \int_{t_n}^t \left[\frac{1}{2} \text{Trace}(\sigma(x)\sigma^*(x) \text{Hess}_x u(s, x)) + \langle \mu(x), \nabla_x u(s, x) \rangle_{\mathbb{R}^d} \right] ds. \end{aligned}$$

To derive the splitting-up approximation, we make a few simplifying assumptions, not all of which are needed to implement the resulting algorithm. First, we suppose that φ has an at most polynomially growing gradient $\nabla\varphi: \mathbb{R}^d \rightarrow \mathbb{R}^d$. Moreover, we assume that there exists a function $v: [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ satisfying $v(0, x) = \varphi(x)$, $x \in \mathbb{R}^d$, such that for every $n \in \{0, 1, \dots, N-1\}$, $v|_{(t_n, t_{n+1}] \times \mathbb{R}^d}$ belongs to $C^{1,2}((t_n, t_{n+1}] \times \mathbb{R}^d, \mathbb{R})$ with at most polynomially growing partial derivatives and for all $n \in \{0, 1, \dots, N-1\}$, $t \in (t_n, t_{n+1}]$, and $x \in \mathbb{R}^d$,

$$(2.7) \quad \begin{aligned} v(t, x) &= v(t_n, x) + f(x, v(t_n, x), \nabla_x v(t_n, x)) (t_{n+1} - t_n) \\ &\quad + \int_{t_n}^t \left[\frac{1}{2} \text{Trace}(\sigma(x)\sigma^*(x) \text{Hess}_x v(s, x)) + \langle \mu(x), \nabla_x v(s, x) \rangle_{\mathbb{R}^d} \right] ds; \end{aligned}$$

see, e.g., Hairer, Hutzenthaler, and Jentzen [46, section 4.4], Deck and Kruse [23], Krylov [65, Chapter 8], and Krylov [66, Theorem 4.32] for existence, uniqueness, and

regularity results for equations of the form (2.7). Comparing (2.7) to (2.6) suggests that

$$(2.8) \quad v(t_n, x) \approx u(t_n, x) \quad \text{for all } n \in \{1, \dots, N\}.$$

So v is a specific splitting-up type approximation of the function u ; see, e.g., [25, 42, 44, 45].

2.2. An approximate Feynman–Kac representation. In the next step, we derive a Feynman–Kac representation of v ; see, e.g., Milstein and Tretyakov [84, section 2]. Let $B: [0, T] \times \Omega \rightarrow \mathbb{R}^d$ be a standard $(\mathcal{F}_t)_{t \in [0, T]}$ -Brownian motion on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in [0, T]}, \mathbb{P})$ satisfying the usual conditions. Consider an $\mathcal{F}_0/\mathcal{B}(\mathbb{R}^d)$ -measurable random variable $\xi: \Omega \rightarrow \mathbb{R}^d$ satisfying $\mathbb{E}[\|\xi\|_{\mathbb{R}^d}^p] < \infty$ for every $p \in (0, \infty)$, and let $Y: [0, T] \times \Omega \rightarrow \mathbb{R}^d$ be an $(\mathcal{F}_t)_{t \in [0, T]}$ -adapted process with continuous sample paths satisfying for every $t \in [0, T]$,

$$(2.9) \quad Y_t = \xi + \int_0^t \mu(Y_s) ds + \int_0^t \sigma(Y_s) dB_s \quad \mathbb{P}\text{-a.s.}$$

The assumption that $\mathbb{E}[\|\xi\|_{\mathbb{R}^d}^p] < \infty$ for all $p \in (0, \infty)$ ensures that

$$(2.10) \quad \sup_{t \in [0, T]} \mathbb{E}[\|Y_t\|_{\mathbb{R}^d}^p] < \infty \quad \text{for all } p \in (0, \infty);$$

see, e.g., Stroock [92, section 1.2]. Moreover, (2.7) implies that for all $n \in \{0, 1, \dots, N-1\}$, $t \in (t_n, t_{n+1})$, and $x \in \mathbb{R}^d$, one has

$$(2.11) \quad \frac{\partial}{\partial t} v(t, x) = \langle \mu(x), \nabla_x v(t, x) \rangle_{\mathbb{R}^d} + \frac{1}{2} \text{Trace}(\sigma(x) \sigma^*(x) \text{Hess}_x v(t, x)),$$

from which it follows that for all $n \in \{0, 1, \dots, N-1\}$, $t \in (T - t_{n+1}, T - t_n)$, and $x \in \mathbb{R}^d$,

$$(2.12) \quad \frac{\partial}{\partial t} v(T - t, x) + \langle \mu(x), \nabla_x v(T - t, x) \rangle_{\mathbb{R}^d} + \frac{1}{2} \text{Trace}(\sigma(x) \sigma^*(x) \text{Hess}_x v(T - t, x)) = 0.$$

Since for every $n \in \{0, 1, \dots, N-1\}$, $v|_{(t_n, t_{n+1}] \times \mathbb{R}^d}$ is in $C^{1,2}((t_n, t_{n+1}] \times \mathbb{R}^d, \mathbb{R})$, we obtain from Itô's formula that for all $n \in \{0, 1, \dots, N-1\}$ and $t \in [T - t_{n+1}, T - t_n)$, we have

$$(2.13) \quad \begin{aligned} v(T - t, Y_t) &= v(T - t_{n+1}, Y_{t_{n+1}}) + \int_{T-t_{n+1}}^t \langle \nabla_x v(T - s, Y_s), \sigma(Y_s) dB_s \rangle_{\mathbb{R}^d} \\ &+ \int_{T-t_{n+1}}^t \frac{\partial}{\partial s} v(T - s, Y_s) ds + \int_{T-t_{n+1}}^t \frac{1}{2} \text{Trace}(\sigma(Y_s) \sigma^*(Y_s) \text{Hess}_x v(T - s, Y_s)) ds \\ &+ \int_{T-t_{n+1}}^t \langle \mu(Y_s), \nabla_x v(T - s, Y_s) \rangle_{\mathbb{R}^d} ds \quad \mathbb{P}\text{-a.s.}, \end{aligned}$$

which by (2.12), gives that $\mathbb{P}$ -a.s.,

$$(2.14) \quad v(T - t, Y_t) = v(T - t_{n+1}, Y_{t_{n+1}}) + \int_{T-t_{n+1}}^t \langle \nabla_x v(T - s, Y_s), \sigma(Y_s) dB_s \rangle_{\mathbb{R}^d}.$$

Moreover, since, by assumption, $\sigma: \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ is Lipschitz continuous and $\nabla_x v(t, x)$ at most polynomially growing in $(t, x) \in (t_n, t_{n+1}] \times \mathbb{R}^d$, one obtains from (2.10) that for all $n \in \{0, 1, \dots, N-1\}$ and $t \in [T - t_{n+1}, T - t_n]$,

$$(2.15) \quad \int_{T-t_{n+1}}^t \mathbb{E} \left[\|\sigma^*(Y_s) \nabla_x v(T-s, Y_s)\|_{\mathbb{R}^d}^2 \right] ds < \infty,$$

from which it follows that

$$(2.16) \quad \mathbb{E} \left[\int_{T-t_{n+1}}^t \langle \nabla_x v(T-s, Y_s), \sigma(Y_s) dB_s \rangle_{\mathbb{R}^d} \middle| \mathcal{F}_{T-t_{n+1}} \right] = 0 \quad \mathbb{P}\text{-a.s.}$$

Together with (2.14), this shows that for all $n \in \{0, 1, \dots, N-1\}$ and $t \in [T - t_{n+1}, T - t_n]$, we have

$$(2.17) \quad \mathbb{E} [v(T-t, Y_t) \mid \mathcal{F}_{T-t_{n+1}}] = \mathbb{E} [v(t_{n+1}, Y_{T-t_{n+1}}) \mid \mathcal{F}_{T-t_{n+1}}] \quad \mathbb{P}\text{-a.s.}$$

Since $Y_{T-t_{n+1}}$ is $\mathcal{F}_{T-t_{n+1}}/\mathcal{B}(\mathbb{R})$ -measurable, one obtains from the tower property of conditional expectations that for all $n \in \{0, 1, \dots, N-1\}$ and $t \in [T - t_{n+1}, T - t_n]$,

$$(2.18) \quad \mathbb{E} [v(T-t, Y_t) \mid Y_{T-t_{n+1}}] = \mathbb{E} [v(t_{n+1}, Y_{T-t_{n+1}}) \mid Y_{T-t_{n+1}}] = v(t_{n+1}, Y_{T-t_{n+1}}) \quad \mathbb{P}\text{-a.s.}$$

Since $(Y_t)_{t \in [0, T]}$ has continuous sample paths and $(v(t, x))_{(t, x) \in (t_n, t_{n+1}] \times \mathbb{R}^d}$ at most polynomially growing first order partial derivatives, it follows from (2.7) and (2.11) that we have for all $n \in \{0, 1, \dots, N-1\}$ and $\omega \in \Omega$,

$$(2.19) \quad \lim_{t \nearrow T-t_n} v(T-t, Y_t(\omega)) = v(t_n, Y_{T-t_n}(\omega)) \\ + f(Y_{T-t_n}(\omega), v(t_n, Y_{T-t_n}(\omega)), \nabla_x v(t_n, Y_{T-t_n}(\omega))) (t_{n+1} - t_n).$$

In addition, since $\sup_{t \in [0, T]} \mathbb{E} [\|Y_t\|_{\mathbb{R}^d}^p] < \infty$ for every $p \in (0, \infty)$, we obtain

$$(2.20) \quad \sup_{t \in (t_n, t_{n+1}]} \mathbb{E} [|v(T-t, Y_t)|^p] < \infty.$$

So it follows from (2.18) and (2.19) that for all $n \in \{0, 1, \dots, N-1\}$,

$$(2.21) \quad v(t_{n+1}, Y_{T-t_{n+1}}) = \mathbb{E} \left[\lim_{t \nearrow T-t_n} v(T-t, Y_t) \mid Y_{T-t_{n+1}} \right] \\ = \mathbb{E} \left[v(t_n, Y_{T-t_n}) + f(Y_{T-t_n}, v(t_n, Y_{T-t_n}), \nabla_x v(t_n, Y_{T-t_n})) (t_{n+1} - t_n) \mid Y_{T-t_{n+1}} \right] \quad \mathbb{P}\text{-a.s.},$$

which is the Feynman–Kac-type representation for which we aim. Note that the nonlinearity f as well as the coefficient functions μ and σ in (2.1) do not depend on t . But the above derivation could be extended to time-dependent f , μ , and σ since the Feynman–Kac formula still holds in this case.

2.3. Formulation as recursive minimization problems. We now reformulate (2.21) as recursive minimization problems. It follows from our assumptions that for every $n \in \{1, 2, \dots, N\}$, $v(t_{n-1}, x) + f(x, v(t_{n-1}, x), \nabla_x v(t_{n-1}, x))(t_n - t_{n-1})$ is at most polynomially growing in $x \in \mathbb{R}^d$. Therefore, we obtain from (2.10) that

$$(2.22) \quad \mathbb{E} \left[|v(t_{n-1}, Y_{T-t_{n-1}}) + f(Y_{T-t_{n-1}}, v(t_{n-1}, Y_{T-t_{n-1}}), \nabla_x v(t_{n-1}, Y_{T-t_{n-1}})) (t_n - t_{n-1})|^2 \right] < \infty.$$

Since $v(t_n, x)$ is continuous in $x \in \mathbb{R}^d$, it follows from the factorization lemma and the L^2 -minimality property of conditional expectations (see, e.g., Klenke [63, Corollary 8.17]) that for every $n \in \{1, 2, \dots, N\}$ we have

(2.23)

$$(v(t_n, x))_{x \in \text{supp}(Y_{T-t_n}(\mathbb{P}))} = \underset{w \in C(\text{supp}(Y_{T-t_n}(\mathbb{P})), \mathbb{R})}{\text{argmin}} \mathbb{E} \left[|w(Y_{T-t_n}) - [v(t_{n-1}, Y_{T-t_{n-1}}) + f(Y_{T-t_{n-1}}, v(t_{n-1}, Y_{T-t_{n-1}}), \nabla_x v(t_{n-1}, Y_{T-t_{n-1}})) (t_n - t_{n-1})]|^2 \right].$$

2.4. Deep artificial neural network approximations. To tackle the minimization problems (2.23) numerically, we approximate the functions $v(t_n, \cdot)$, $n \in \{1, 2, \dots, N\}$, with neural networks V_n . More precisely, we choose $k \in \{3, 4, \dots\}$ and $l \in \mathbb{N}$. Then we set $\nu = (1 + kl - l)(l + 1) + l(d + 1)$ and consider functions $V_n: \mathbb{R}^\nu \times \mathbb{R}^d \rightarrow \mathbb{R}$, $n \in \{0, 1, \dots, N\}$, such that for every $(\theta, x) \in \mathbb{R}^\nu \times \mathbb{R}^d$, V_0 is given by

$$(2.24) \quad V_0(\theta, x) = \varphi(x)$$

and for $n \in \{1, 2, \dots, N\}$, V_n is of the form

(2.25)

$$V_n(\theta, x) = A_{l,1}^{\theta, (k-1)l(l+1)+l(d+1)} \circ \mathcal{L}_l \circ A_{l,l}^{\theta, (k-2)l(l+1)+l(d+1)} \circ \dots \circ \mathcal{L}_l \circ A_{l,l}^{\theta, l(d+1)} \circ \mathcal{L}_l \circ A_{d,l}^{\theta, 0}(x),$$

where for $r \in \mathbb{N}_0 = \{0\} \cup \mathbb{N}$ and $i, j \in \mathbb{N}$ with $r + j(i + 1) \leq \nu$, $A_{i,j}^{\theta, r}: \mathbb{R}^i \rightarrow \mathbb{R}^j$ is the affine function defined by

(2.26)

$$A_{i,j}^{\theta, r}(x) = \begin{pmatrix} \theta_{r+1} & \theta_{r+2} & \dots & \theta_{r+i} \\ \theta_{r+i+1} & \theta_{r+i+2} & \dots & \theta_{r+2i} \\ \theta_{r+2i+1} & \theta_{r+2i+2} & \dots & \theta_{r+3i} \\ \vdots & \vdots & \vdots & \vdots \\ \theta_{r+i(j-1)+1} & \theta_{r+i(j-1)+2} & \dots & \theta_{r+ij} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_i \end{pmatrix} + \begin{pmatrix} \theta_{r+ij+1} \\ \theta_{r+ij+2} \\ \theta_{r+ij+3} \\ \vdots \\ \theta_{r+ij+j} \end{pmatrix},$$

and $\mathcal{L}_l: \mathbb{R}^l \rightarrow \mathbb{R}^l$ is a mapping of the form

$$(2.27) \quad \mathcal{L}_l(x_1, \dots, x_l) = (\rho(x_1), \dots, \rho(x_l))$$

for a weakly differentiable continuous function $\rho: \mathbb{R} \rightarrow \mathbb{R}$.

Equation (2.25) is a feedforward neural network with activation function ρ and $k + 1$ layers (an input layer with d neurons, $k - 1$ hidden layers with l neurons each, and an output layer with one neuron); see, e.g., [9, 70]. Commonly used activation functions are, e.g., the logistic function $x \mapsto e^x / (1 + e^x)$ or the ReLU function $x \mapsto \max\{0, x\}$. The logistic function is continuously differentiable. So the corresponding neural networks V_n have well-defined θ - and x -gradients $\nabla_\theta V_n$ and $\nabla_x V_n$. However, in the examples of section 3 below we use the ReLU function. It is continuously differentiable on $(-\infty, 0) \cup (0, \infty)$ and has a left-hand derivative at 0. This yields weak θ - and x -gradients for V_n , which we also denote by $\nabla_\theta V_n$ and $\nabla_x V_n$, respectively.

2.5. Stochastic gradient descent based minimization. Next, we recursively solve quadratic minimization problems to find parameter vectors $\theta^1, \theta^2, \dots, \theta^N \in \mathbb{R}^\nu$ such that $V_n(\theta^n, x) \approx v(t_n, x)$ for $n \in \{1, 2, \dots, N\}$. More specifically, $\theta^0 \in \mathbb{R}^\nu$ can be chosen arbitrarily; e.g., $\theta^0 = (0, \dots, 0) \in \mathbb{R}^\nu$. For given $n \in \{1, 2, \dots, N\}$ and

$\theta^0, \theta^1, \dots, \theta^{n-1} \in \mathbb{R}^\nu$, one tries to find an approximate minimizer $\theta^n \in \mathbb{R}^\nu$ of the function

$$(2.28) \quad \mathbb{R}^\nu \ni \theta \mapsto \mathbb{E} \left[\left| V_n(\theta, Y_{T-t_n}) - [V_{n-1}(\theta^{n-1}, Y_{T-t_{n-1}}) + f(Y_{T-t_{n-1}}, V_{n-1}(\theta^{n-1}, Y_{T-t_{n-1}}), \nabla_x V_{n-1}(\theta^{n-1}, Y_{T-t_{n-1}})) (t_n - t_{n-1})] \right|^2 \right] \in \mathbb{R}.$$

To do this with a standard stochastic gradient descent, one can initialize ϑ_0^n randomly or, e.g., as $\vartheta_0^n = \theta^{n-1}$. Then one chooses a stepsize $\gamma \in (0, \infty)$ together with a number $M \in \mathbb{N}$ and iteratively updates for $m \in \{0, 1, \dots, M-1\}$ according to

$$(2.29) \quad \vartheta_{m+1}^n = \vartheta_m^n - 2\gamma \nabla_\theta V_n(\vartheta_m^n, Y_{T-t_n}^m) \left[V_n(\vartheta_m^n, Y_{T-t_n}^m) - V_{n-1}(\theta^{n-1}, Y_{T-t_{n-1}}^m) - f(Y_{T-t_{n-1}}^m, V_{n-1}(\theta^{n-1}, Y_{T-t_{n-1}}^m), \nabla_x V_{n-1}(\theta^{n-1}, Y_{T-t_{n-1}}^m)) (t_n - t_{n-1}) \right],$$

where $Y^m: [0, T] \times \Omega \rightarrow \mathbb{R}^d$, $m \in \{0, 1, \dots, M-1\}$, are $(\mathcal{F}_t)_{t \in [0, T]}$ -adapted stochastic processes with continuous sample paths satisfying for every $t \in [0, T]$ the SDEs

$$(2.30) \quad Y_t^m = \xi^m + \int_0^t \mu(Y_s^m) ds + \int_0^t \sigma(Y_s^m) dB_s^m \quad \mathbb{P}\text{-a.s.}$$

corresponding to i.i.d. standard $(\mathcal{F}_t)_{t \in [0, T]}$ -Brownian motions $B^m: [0, T] \times \Omega \rightarrow \mathbb{R}^d$ and i.i.d. $\mathcal{F}_0/\mathcal{B}(\mathbb{R}^d)$ -measurable functions $\xi^m: \Omega \rightarrow \mathbb{R}^d$, $m \in \{0, 1, \dots, M-1\}$. After M gradient steps one sets $\theta^n = \vartheta_M^n$.

2.6. Discretization of the auxiliary stochastic process Y . Equation (2.29) provides an implementable numerical algorithm in the special case where the solutions Y^m of the SDEs (2.30) can be simulated exactly. If this is not the case, one can use a numerical approximation method to approximatively simulate Y^m , $m \in \{0, 1, \dots, M-1\}$. In the following, we concentrate on the Euler–Maruyama scheme. But it is also possible to use a different approximation method.

Note that it follows from (2.30) that for all $m \in \{0, 1, \dots, M-1\}$ and $n \in \{0, 1, \dots, N-1\}$ we have

$$(2.31) \quad Y_{T-t_n}^m = Y_{T-t_{n+1}}^m + \int_{T-t_{n+1}}^{T-t_n} \mu(Y_s^m) ds + \int_{T-t_{n+1}}^{T-t_n} \sigma(Y_s^m) dB_s^m \quad \mathbb{P}\text{-a.s.},$$

or equivalently,

$$(2.32) \quad Y_{\tau_{n+1}}^m = Y_{\tau_n}^m + \int_{\tau_n}^{\tau_{n+1}} \mu(Y_s^m) ds + \int_{\tau_n}^{\tau_{n+1}} \sigma(Y_s^m) dB_s^m \quad \mathbb{P}\text{-a.s.}$$

for $\tau_n = T - t_{N-n}$. This suggests that for all $m \in \{0, 1, \dots, M-1\}$ and $n \in \{0, 1, \dots, N-1\}$,

$$(2.33) \quad Y_{\tau_{n+1}}^m \approx Y_{\tau_n}^m + \mu(Y_{\tau_n}^m) (\tau_{n+1} - \tau_n) + \sigma(Y_{\tau_n}^m) (B_{\tau_{n+1}}^m - B_{\tau_n}^m).$$

Therefore, we introduce the Euler–Maruyama approximations $\mathcal{Y}^m: \{0, 1, \dots, N\} \times \Omega \rightarrow \mathbb{R}^d$ given for every $n \in \{0, 1, \dots, N-1\}$ by $\mathcal{Y}_0^m = \xi^m$ and

$$(2.34) \quad \mathcal{Y}_{n+1}^m = \mathcal{Y}_n^m + \mu(\mathcal{Y}_n^m) (\tau_{n+1} - \tau_n) + \sigma(\mathcal{Y}_n^m) (B_{\tau_{n+1}}^m - B_{\tau_n}^m).$$

It can be seen from (2.33) and (2.34) that for every $n \in \{0, 1, \dots, N\}$, one has

$$(2.35) \quad \mathcal{Y}_n^m \approx Y_{\tau_n}^m = Y_{T-t_{N-n}}^m \quad \text{and hence} \quad Y_{T-t_n}^m \approx \mathcal{Y}_{N-n}^m,$$

which can be used to derive approximations of $(\vartheta_m^n)_{m=0}^M$, $n \in \{1, 2, \dots, N\}$, from (2.29) that are also implementable if the processes Y^m cannot be simulated exactly. More precisely, set $\Theta_M^0 = (0, \dots, 0) \in \mathbb{R}^\nu$. For $n \in \{1, 2, \dots, N\}$, initialize Θ_0^n randomly or as $\Theta_0^n = \Theta_M^{n-1}$. Then set for every $m \in \{0, 1, \dots, M-1\}$,

$$(2.36) \quad \Theta_{m+1}^n = \Theta_m^n - 2\gamma \nabla_\theta V_n(\Theta_m^n, \mathcal{Y}_{N-n}^m) \left[V_n(\Theta_m^n, \mathcal{Y}_{N-n}^m) - V_{n-1}(\Theta_M^{n-1}, \mathcal{Y}_{N-n+1}^m) \right. \\ \left. - f(\mathcal{Y}_{N-n+1}^m, V_{n-1}(\Theta_M^{n-1}, \mathcal{Y}_{N-n+1}^m), \nabla_x V_{n-1}(\Theta_M^{n-1}, \mathcal{Y}_{N-n+1}^m)) (t_n - t_{n-1}) \right].$$

Comparing (2.36) to (2.29) suggests that $\Theta_m^n \approx \vartheta_m^n$ for $m \in \{0, 1, \dots, M\}$ and $n \in \{1, 2, \dots, N\}$.

In subsections 2.7 and 2.8, we first formalize and then generalize the approximation algorithm derived in subsections 2.1–2.6.

2.7. Description of the algorithm in a special case. In this subsection, we provide a formal description of the algorithm derived in subsections 2.1–2.6 in the case where the standard Euler–Maruyama scheme (cf., e.g., [64, 79, 81]) is used to approximate the solutions Y^m of the SDEs (2.30) and optimal parameters $\theta_1, \dots, \theta_N \in \mathbb{R}^\nu$ are computed with a plain vanilla stochastic gradient descent with a constant learning rate $\gamma \in (0, \infty)$ and batch size 1. Note that for the algorithm to be implementable, it is enough if the initial condition φ has a weak gradient $\nabla\varphi: \mathbb{R}^d \rightarrow \mathbb{R}^d$ which does not need to satisfy any growth conditions.

In Framework 2.1, feedforward neural networks of the form (2.25) are employed to approximate the solution of the PDE. Equation (2.39) describes a stochastic gradient descent scheme with constant learning rate γ , and (2.40) specifies the quadratic loss functions.

FRAMEWORK 2.1 (special case). Assume φ has a weak gradient $\nabla\varphi: \mathbb{R}^d \rightarrow \mathbb{R}^d$. Consider $N \in \mathbb{N}$ and $t_0, t_1, \dots, t_N \in [0, T]$ such that

$$(2.37) \quad 0 = t_0 < t_1 < \dots < t_N = T.$$

Set $\tau_n = T - t_{N-n}$ for $n \in \{0, 1, \dots, N\}$. For a given $M \in \mathbb{N}$, let $B^m: [0, T] \times \Omega \rightarrow \mathbb{R}^d$, $m \in \{0, 1, \dots, M-1\}$, be i.i.d. standard Brownian motions on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Consider i.i.d. random variables $\xi^m: \Omega \rightarrow \mathbb{R}^d$, $m \in \{0, 1, \dots, M-1\}$, that are independent of B^m , $m \in \{0, 1, \dots, M-1\}$, and let the stochastic processes $\mathcal{Y}^m: \{0, 1, \dots, N\} \times \Omega \rightarrow \mathbb{R}^d$, $m \in \{0, 1, \dots, M-1\}$, be given by $\mathcal{Y}_0^m = \xi^m$ and

$$(2.38) \quad \mathcal{Y}_{n+1}^m = \mathcal{Y}_n^m + \mu(\mathcal{Y}_n^m)(\tau_{n+1} - \tau_n) + \sigma(\mathcal{Y}_n^m)(B_{\tau_{n+1}}^m - B_{\tau_n}^m), \quad n \in \{0, 1, \dots, N-1\}.$$

Let $V_n: \mathbb{R}^\nu \times \mathbb{R}^d \rightarrow \mathbb{R}$, $n \in \{0, 1, \dots, N\}$, be the functions given in (2.24)–(2.25). Consider $\gamma \in (0, \infty)$, and let $\Theta^n: \{0, 1, \dots, M\} \times \Omega \rightarrow \mathbb{R}^\nu$, $n \in \{0, 1, \dots, N\}$, be stochastic processes satisfying for all $m \in \{0, 1, \dots, M-1\}$ and $n \in \{1, \dots, N\}$,

$$(2.39) \quad \Theta_{m+1}^n = \Theta_m^n - \gamma \Phi^{n,m}(\Theta_m^n),$$

where for all $\omega \in \Omega$,

(2.40)

$$\begin{aligned} \phi^{n,m}(\theta, \omega) = & \left[V_n(\theta, \mathcal{Y}_{N-n}^m(\omega)) - V_{n-1}(\Theta_M^{n-1}(\omega), \mathcal{Y}_{N-n+1}^m(\omega)) - (t_n - t_{n-1}) \right. \\ & \left. \times f(\mathcal{Y}_{N-n+1}^m(\omega), V_{n-1}(\Theta_M^{n-1}(\omega), \mathcal{Y}_{N-n+1}^m(\omega)), \nabla_x V_{n-1}(\Theta_M^{n-1}(\omega), \mathcal{Y}_{N-n+1}^m(\omega))) \right]^2, \end{aligned}$$

and $\Phi^{n,m}(\theta, \omega) = \nabla_\theta \phi^{n,m}(\theta, \omega)$.

In the setting of Framework 2.1, the random variables $V_n(\Theta_M^n, x): \Omega \rightarrow \mathbb{R}$ provide for all $n \in \{1, 2, \dots, N\}$ and $x \in \mathbb{R}^d$ the approximations

$$(2.41) \quad V_n(\Theta_M^n, x) \approx u(t_n, x)$$

to the solution $u: [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ of the PDE (2.1).

2.8. Description of the algorithm in the general case. We now generalize Framework 2.1 so that, besides a plain vanilla stochastic gradient descent with a constant learning rate and batch size 1, it also covers more advanced machine learning techniques such as mini-batches, batch normalization, and more sophisticated updating rules.

In Framework 2.2 below, functions $V_n^{j,s}: \mathbb{R}^\nu \times \mathbb{R}^d \rightarrow \mathbb{R}$ parametrized by $(j, s, n) \in \mathbb{N} \times \mathbb{R}^\varsigma \times \{0, 1, \dots, N\}$ for some number $\varsigma \in \mathbb{N}$ are used to approximate the solution of the PDE. The additional parameters j and s make it possible to describe mini-batches and batch normalization; see Ioffe and Szegedy [59]. As in Framework 2.1, the standard Euler–Maruyama scheme is used to approximate the solutions Y^m of the SDEs (2.30). But a different approximation scheme could be employed as well. In (2.43), the quadratic loss functions are given that are used for training the functions $V_n^{j,s}$, whereas (2.44) specifies the gradients of the loss functions. The role of the stochastic processes $\mathbb{S}^n$ in (2.45) is to describe the variables (running mean and standard deviation) needed for batch normalization. The stochastic processes Θ^n and Ξ^n in (2.46) describe the updating rule. Their dynamics are specified by the functions Ψ_m^n and ψ_m^n . Since we make no assumptions on Ψ_m^n and ψ_m^n , the framework includes very general stochastic gradient optimization methods. In the examples in section 3, we use Adam optimization; see Kingma and Ba [62]. The corresponding specification of the functions Ψ_m^n and ψ_m^n is given in (3.1)–(3.2) below.

FRAMEWORK 2.2 (general case). Assume φ has a weak gradient $\nabla\varphi: \mathbb{R}^d \rightarrow \mathbb{R}^d$. Consider $N \in \mathbb{N}$ and $t_0, t_1, \dots, t_N \in [0, T]$ such that $0 = t_0 < t_1 < \dots < t_N = T$. Set $\tau_n = T - t_{N-n}$ for $n \in \{0, 1, \dots, N\}$. For a given $M \in \mathbb{N}$, let, for every $n \in \{1, 2, \dots, N\}$, $B^{n,m,j}: [0, T] \times \Omega \rightarrow \mathbb{R}^d$, $m \in \{0, 1, \dots, M-1\}$, $j \in \mathbb{N}$, be i.i.d. standard Brownian motions on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and $\xi^{n,m,j}: \Omega \rightarrow \mathbb{R}^d$, $m \in \{0, 1, \dots, M-1\}$, $j \in \mathbb{N}$, i.i.d. random variables that are independent of $B^{n,m,j}$, $m \in \{0, 1, \dots, M-1\}$, $j \in \mathbb{N}$. Let $\mathcal{Y}^{n,m,j}: \{0, 1, \dots, N\} \times \Omega \rightarrow \mathbb{R}^d$, $n \in \{1, 2, \dots, N\}$, $m \in \{0, 1, \dots, M-1\}$, $j \in \mathbb{N}$, be stochastic processes given by $\mathcal{Y}_0^{n,m,j} = \xi^{n,m,j}$ and

$$(2.42) \quad \mathcal{Y}_{k+1}^{n,m,j} = \mu(\mathcal{Y}_k^{n,m,j})(\tau_{k+1} - \tau_k) + \sigma(\mathcal{Y}_k^{n,m,j})(B_{\tau_{k+1}}^{n,m,j} - B_{\tau_k}^{n,m,j}), \quad k \in \{0, 1, \dots, N-1\}.$$

Let $\nu, \varsigma, \varrho, J_0, \dots, J_{M-1} \in \mathbb{N}$. Consider functions $V_n^{j,s}: \mathbb{R}^\nu \times \mathbb{R}^d \rightarrow \mathbb{R}$, $(j, s, n) \in \mathbb{N} \times \mathbb{R}^\varsigma \times \{0, 1, \dots, N\}$, such that $V_0^{j,s}(\theta, x) = \varphi(x)$ for all $(j, s, \theta, x) \in \mathbb{N} \times \mathbb{R}^\varsigma \times \mathbb{R}^\nu \times \mathbb{R}^d$, and let $\Theta^n: \{0, 1, \dots, M-1\} \times \Omega \rightarrow \mathbb{R}^\nu$, $n \in \{0, 1, \dots, N\}$, be stochastic processes. For all $n \in \{1, 2, \dots, N\}$, $m \in \{0, 1, \dots, M-1\}$, and $s \in \mathbb{R}^\varsigma$, let the

mapping $\phi^{n,m,s}: \mathbb{R}^\nu \times \Omega \rightarrow \mathbb{R}$ be given by

(2.43)

$$\begin{aligned} \phi^{n,m,s}(\theta, \omega) = & \frac{1}{J_m} \sum_{j=1}^{J_m} \left[V_n^{j,s}(\theta, \mathcal{Y}_{N-n}^{n,m,j}(\omega)) - V_{n-1}^{j,s}(\Theta_M^{n-1}(\omega), \mathcal{Y}_{N-n+1}^{n,m,j}(\omega)) - (t_n - t_{n-1}) \right. \\ & \left. \times f\left(\mathcal{Y}_{N-n+1}^{n,m,j}(\omega), V_{n-1}^{j,s}(\Theta_M^{n-1}(\omega), \mathcal{Y}_{N-n+1}^{n,m,j}(\omega)), \nabla_x V_{n-1}^{j,s}(\Theta_M^{n-1}(\omega), \mathcal{Y}_{N-n+1}^{n,m,j}(\omega))\right) \right]^2, \end{aligned}$$

and assume $\Phi^{n,m,s}: \mathbb{R}^\nu \times \Omega \rightarrow \mathbb{R}^\nu$ is a function satisfying

$$(2.44) \quad \Phi^{n,m,s}(\theta, \omega) = \nabla_\theta \phi^{n,m,s}(\theta, \omega)$$

for all $\omega \in \Omega$ and $\theta \in \{\vartheta \in \mathbb{R}^\nu: \phi^{n,m,s}(\cdot, \omega): \mathbb{R}^\nu \rightarrow \mathbb{R} \text{ is differentiable at } \vartheta\}$. For every $n \in \{1, 2, \dots, N\}$ and $m \in \{0, 1, \dots, M-1\}$, let $\mathcal{S}^n: \mathbb{R}^s \times \mathbb{R}^\nu \times (\mathbb{R}^d)^{\{0,1,\dots,N\} \times \mathbb{N}} \rightarrow \mathbb{R}^s$, $\Psi_m^n: \mathbb{R}^e \times \mathbb{R}^\nu \rightarrow \mathbb{R}^e$, and $\psi_m^n: \mathbb{R}^e \rightarrow \mathbb{R}^\nu$ be functions and $\mathbb{S}^n: \{0, 1, \dots, M-1\} \times \Omega \rightarrow \mathbb{R}^s$ and $\Xi^n: \{0, 1, \dots, M-1\} \times \Omega \rightarrow \mathbb{R}^e$ stochastic processes such that

$$(2.45) \quad \mathbb{S}_{m+1}^n = \mathcal{S}^n(\mathbb{S}_m^n, \Theta_m^n, (\mathcal{Y}_k^{n,m,i})_{(k,i) \in \{0,1,\dots,N\} \times \mathbb{N}}),$$

$$(2.46) \quad \Xi_{m+1}^n = \Psi_m^n(\Xi_m^n, \Phi^{n,m,\mathbb{S}_{m+1}^n}(\Theta_m^n)) \quad \text{and} \quad \Theta_{m+1}^n = \Theta_m^n - \psi_m^n(\Xi_{m+1}^n).$$

In the setting of Framework 2.2, the functions $V_n^{1,\mathbb{S}_M^n}(\Theta_M^n, x): \Omega \rightarrow \mathbb{R}$ yield the approximations

$$(2.47) \quad V_n^{1,\mathbb{S}_M^n}(\Theta_M^n, x) \approx u(t_n, x), \quad n \in \{1, 2, \dots, N\}, \quad x \in \mathbb{R}^d,$$

of the solution $u: [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ to the PDE (2.1).

3. Examples. We now illustrate the performance of the deep splitting method on five concrete example PDEs. In each example, we use the general approximation method of Framework 2.2 with approximating functions $V_n^{j,s}: \mathbb{R}^\nu \times \mathbb{R}^d \rightarrow \mathbb{R}$, $n \in \{1, 2, \dots, N\}$, specified as feedforward neural networks with 4 layers (1 input layer, 2 hidden layers, and 1 output layer) and ReLU-activation $\rho(x) = \max\{0, x\}$, $x \in \mathbb{R}$. We use mini-batches of size $J_m = 256$ and apply batch normalization before the first affine transformation, before each of the two nonlinear activation functions in front of the hidden layers, and just before the output layer. We use Xavier initialization (see Glorot and Bengio [34]) to initialize all weights in the neural networks together with Adam optimization (see Kingma and Ba [62]) with parameters $\varepsilon = 10^{-8}$, $\beta_1 = 0.9$, $\beta_2 = 0.999$ and decreasing learning rates $(\gamma_m)_{m=0}^{M-1}$ that we choose depending on the form and dimension of the problem. More precisely, we set $\varrho = 2\nu$ and denote by $\text{Pow}_2: \mathbb{R}^\nu \rightarrow \mathbb{R}^\nu$ the function given by $\text{Pow}_2(\eta_1, \dots, \eta_\nu) = (\eta_1^2, \dots, \eta_\nu^2)$. Then Adam optimization corresponds to the following specification of the two functions $\Psi_m^n: \mathbb{R}^{3\nu} \rightarrow \mathbb{R}^{2\nu}$ and $\psi_m^n: \mathbb{R}^{2\nu} \rightarrow \mathbb{R}^\nu$ from Framework 2.2:

$$(3.1) \quad \Psi_m^n(x, y, \eta) = (\beta_1 x + (1 - \beta_1)\eta, \beta_2 y + (1 - \beta_2)\text{Pow}_2(\eta))$$

and

$$(3.2) \quad \psi_m^n(x, y) = \left(\left[\sqrt{\frac{|y_1|}{1 - \beta_2^m}} + \varepsilon \right]^{-1} \frac{\gamma_m x_1}{1 - \beta_1^m}, \dots, \left[\sqrt{\frac{|y_\nu|}{1 - \beta_2^m}} + \varepsilon \right]^{-1} \frac{\gamma_m x_\nu}{1 - \beta_1^m} \right).$$

In the examples in the following subsections, we approximate $u(T, x)$ for different $T \in (0, \infty)$ and $x \in \mathbb{R}^d$ with $V_N^{1, S_M^N}(\Theta_M^N, x)$, which, due to the stochastic gradient optimization method, is a random variable. In each example, we report estimates of the expectation and standard deviation of $V_N^{1, S_M^N}(\Theta_M^N, x)$. We also give relative L^1 -approximation errors with respect to reference values calculated with different alternative methods together with their uncorrected sample standard deviations. The average runtimes needed for calculating one realization of $V_N^{1, S_M^N}(\Theta_M^N, x)$ are in seconds and were determined as averages over 10 independent runs.

All numerical experiments presented in this paper were implemented in PYTHON using TENSORFLOW and run on a NVIDIA GeForce GTX 1080 GPU with 1974 MHz core clock and 8 GB GDDR5X memory with 1809.5 MHz clock rate and an underlying system consisting of an Intel Core i7-6800K 3.4 GHz CPU with 64 GB DDR4-2133 memory running TensorFlow 1.5 on Ubuntu 16.04. The PYTHON source codes are available online from https://github.com/seb-becker/deep_pde.

3.1. Hamilton–Jacobi–Bellman (HJB) equations. As a first example, we calculate approximate solutions of the PDE

$$(3.3) \quad \frac{\partial}{\partial t} u(t, x) = \Delta_x u(t, x) - \|\nabla_x u(t, x)\|_{\mathbb{R}^d}^2, \quad (t, x) \in (0, T] \times \mathbb{R}^d,$$

with initial condition $u(0, x) = \|x\|_{\mathbb{R}^d}^{1/2}$ for different $d \in \mathbb{N}$. The deep splitting method can be applied to more general HJB equations. But (3.3) has the advantage that it reduces to a linear heat equation under a logarithmic transformation; see, e.g., E, Han, and Jentzen [27, Lemma 4.2]. Since solutions of linear equations can be approximated with standard Monte Carlo, this allows us to efficiently compute reference solutions in high dimensions.

Table 1 shows deep splitting approximations of $u(T, 0, \dots, 0)$ for different values of d , T , and N . We derived it with $\mu(x) = (0, 0, \dots, 0) \in \mathbb{R}^d$, $\sigma(x) = \sqrt{2} \text{Id}_{\mathbb{R}^d \times d}$, and $f(x, y, z) = -\|z\|_{\mathbb{R}^d}^2$ for $(x, y, z) \in \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d$. We used $M = 500 + 100 \mathbb{1}_{\{10,000\}}(d)$ and

$$(3.4) \quad \gamma_m = \begin{cases} 10^{-1} \mathbb{1}_{[0,300]}(m) + 10^{-2} \mathbb{1}_{(300,400]}(m) + 10^{-3} \mathbb{1}_{(400,500]}(m) & \text{for } d < 10,000, \\ 10^{-1} \mathbb{1}_{[0,400]}(m) + 10^{-2} \mathbb{1}_{(400,500]}(m) + 10^{-3} \mathbb{1}_{(500,600]}(m) & \text{for } d = 10,000. \end{cases}$$

We set $\xi^{n,m,j} = (0, 0, \dots, 0) \in \mathbb{R}^d$ for every $(n, m, j) \in \{1, 2, \dots, N\} \times \{0, 1, \dots, M-1\} \times \mathbb{N}$ and used feedforward neural networks with a d -dimensional input layer, two hidden layers of dimension $d+10$, and a one-dimensional output layer.

The reference values for $u(T, 0, 0, \dots, 0)$ were calculated using a logarithmic transformation and a standard Monte Carlo method; see Han, Jentzen, and E [47, Lemma 4.2].

3.2. Nonlinear Black–Scholes equations. There exist a number of extensions of the classical linear Black–Scholes equation which incorporate nonlinear phenomena such as transaction costs, default risk, or Knightian uncertainty. We here consider nonlinear Black–Scholes equations of the form

$$(3.5) \quad \begin{aligned} \frac{\partial}{\partial t} u(t, x) = & -u(t, x) (1 - \delta) \left[\min \left\{ \gamma^h, \max \left\{ \gamma^l, \frac{(\gamma^h - \gamma^l)}{(v^h - v^l)} (u(t, x) - v^h) + \gamma^h \right\} \right\} \right] \\ & - Ru(t, x) + \langle \bar{\mu} x, \nabla_x u(t, x) \rangle_{\mathbb{R}^d} + \frac{\bar{\sigma}^2}{2} \left[\sum_{i=1}^d |x_i|^2 \frac{\partial^2}{\partial x_i^2} u(t, x) \right], \end{aligned}$$

TABLE 1

Deep splitting approximations of the solution of the HJB equation (3.3) for different values of d , T and N .

d	T	N	Expectation	Std. dev.	Ref. value	Rel. L^1 -error	Std. dev. rel. error	Avg. runtime
10	1/3	8	1.56645	0.00246699	1.56006	0.00410	0.00158134	18.0s
10	2/3	16	1.86402	0.00338646	1.85150	0.00677	0.00182904	37.9s
10	1	24	2.07017	0.00634850	2.04629	0.01167	0.00310245	58.2s
50	1/3	8	2.39214	0.00151918	2.38654	0.00234	0.00063656	18.0s
50	2/3	16	2.84607	0.00140300	2.83647	0.00338	0.00049463	37.9s
50	1	24	3.15098	0.00275839	3.13788	0.00417	0.00087906	58.4s
100	1/3	8	2.85090	0.00071267	2.84696	0.00138	0.00025033	18.1s
100	2/3	16	3.39109	0.00093368	3.38450	0.00195	0.00027587	38.2s
100	1	24	3.75329	0.00136920	3.74471	0.00229	0.00036564	58.3s
200	1/3	8	3.39423	0.00051028	3.39129	0.00087	0.00015047	18.1s
200	2/3	16	4.03680	0.00088215	4.03217	0.00115	0.00021878	38.0s
200	1	24	4.46734	0.00079688	4.46172	0.00126	0.00017860	58.2s
300	1/3	8	3.75741	0.00063334	3.75530	0.00056	0.00016865	18.3s
300	2/3	16	4.46859	0.00049953	4.46514	0.00077	0.00011187	38.5s
300	1	24	4.94586	0.00087736	4.94105	0.00097	0.00017756	58.8s
500	1/3	8	4.27079	0.00051256	4.26900	0.00042	0.00012007	18.0s
500	2/3	16	5.07900	0.00034792	5.07618	0.00056	0.00006854	38.0s
500	1	24	5.62126	0.00045092	5.61735	0.00070	0.00008027	57.7s
1000	1/3	8	5.07989	0.00022764	5.07876	0.00022	0.00004482	20.6s
1000	2/3	16	6.04130	0.00030680	6.03933	0.00033	0.00005080	43.7s
1000	1	24	6.68594	0.00040334	6.68335	0.00039	0.00006035	66.5s
5,000	1/3	8	7.59772	0.00024745	7.59733	0.00005	0.00003257	120.4s
5,000	2/3	16	9.03721	0.00027322	9.03466	0.00028	0.00003024	256.7s
5,000	1	24	9.97266	0.00047098	9.99835	0.00257	0.00004711	393.9s
10,000	1/3	8	9.03574	0.00022994	9.03535	0.00004	0.00002545	519.8s
10,000	2/3	16	10.74521	0.00026228	10.74478	0.00004	0.00002157	1105.6s
10,000	1	24	11.87860	0.00022705	11.89099	0.00104	0.00001909	1687.7s

TABLE 2

Deep splitting approximations of the solution of the nonlinear Black–Scholes equations (3.5) for $T = 1/3$, $N = 96$, and different d .

d	Expectation	Std. dev.	Ref. value	Rel. L^1 -error	Std. dev. rel. error	Avg. runtime
10	40.6553107	0.1000347132	40.7611353	0.0029624273	0.0019393471	858.3s
50	37.421057	0.0339765334	37.5217732	0.0026842068	0.0009055151	975.4s
100	36.3498646	0.027989905	36.4084035	0.0016078403	0.000768776	1481.5s
200	35.374638	0.035236816	35.4127342	0.0012857962	0.0006625744	951.2s
300	34.8476466	0.0225350305	34.8747946	0.0008818254	0.0004762554	953.3s
500	34.2206181	0.0081072294	34.2357988	0.0004701552	0.0001701012	956.0s
1000	33.4058827	0.0050161752	33.4358163	0.0008952555	0.000150024	1039.6s
5,000	31.7511529	0.0048508218	31.7906594	0.0012427078	0.0001525864	7229.7s
10,000	31.1215014	0.0031131196	31.1569116	0.0011365119	0.00009991746	23,593.2s

$(t, x) \in (0, T] \times \mathbb{R}^d$, for suitable parameters $\delta, R, \gamma^h, \gamma^l, v^h, v^l, \bar{\mu}, \bar{\sigma} \in \mathbb{R}$. They describe derivative prices under default risk; see, e.g., Han, Jentzen, and E [47, subsection 3.1] and E et al. [29, subsection 3.1], from where we adopt the initial condition $u(0, x) = \min_{i \in \{1, 2, \dots, d\}} x_i$ and the choice of the parameter values $\delta = 2/3$, $R = 0.02$, $\gamma^h = 0.2$, $\gamma^l = 0.02$, $v^h = 50$, $v^l = 70$, $\bar{\mu} = 0.02$, $\bar{\sigma} = 0.2$.

Table 2 reports deep splitting approximations of $u(T, 50, \dots, 50)$ for $T = 1/3$, $N = 96$, and different values of d . We chose $\mu(x) = \bar{\mu}x$, $\sigma(x) = \bar{\sigma}x$, and

$$(3.6) \quad f(x, y, z) = -(1 - \delta) \min \left\{ \gamma^h, \max \left\{ \gamma^l, \frac{(\gamma^h - \gamma^l)}{(v^h - v^l)} (y - v^h) + \gamma^h \right\} \right\} y - Ry,$$

$(x, y, z) \in \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d$, and used $M = 2000 + 1000 \mathbb{1}_{[0, 100]}(d)$ together with

$$(3.7) \quad \gamma_m = \begin{cases} 10^{-1} \mathbb{1}_{[0, 2500]}(m) + 10^{-2} \mathbb{1}_{(2500, 2750]}(m) + 10^{-3} \mathbb{1}_{(2750, 3000]}(m) & \text{for } d \leq 100, \\ 10^{-1} \mathbb{1}_{[0, 1500]}(m) + 10^{-2} \mathbb{1}_{(1500, 1750]}(m) + 10^{-3} \mathbb{1}_{(1750, 2000]}(m) & \text{for } d > 100. \end{cases}$$

TABLE 3

Deep splitting approximations of solutions of Allen–Cahn-type equations of the form (3.8) for $T = 0.3$, $N = 10$, and different d .

d	Expectation	Std. dev.	Ref. value	Rel. L^1 -error	Std. dev. rel. error	Avg. runtime
10	0.89327	0.00299962	0.89060	0.00364	0.00258004	22.7s
50	1.01855	0.00073173	1.01830	0.00063	0.00036976	22.2s
100	1.04348	0.00029431	1.04510	0.00156	0.00028161	22.3s
200	1.06119	0.00018821	1.06220	0.00096	0.00017719	22.4s
300	1.06961	0.00017250	1.07217	0.00239	0.00016089	22.6s
500	1.07847	0.00013055	1.08124	0.00256	0.00012074	23.1s
1000	1.08842	0.00005689	1.09100	0.00236	0.00005215	25.9s
5,000	1.10522	0.00005201	1.10691	0.00153	0.00004699	134.6s
10,000	1.11071	0.00004502	1.11402	0.00296	0.00004041	473.6s

We set $\xi^{n,m,j} = (50, 50, \dots, 50) \in \mathbb{R}^d$ for every $(n, m, j) \in \{1, 2, \dots, N\} \times \{0, 1, \dots, M-1\} \times \mathbb{N}$ and used feedforward neural networks with a d -dimensional input layer, two hidden layers of dimension $d + 10 + 40 \mathbb{1}_{[1,100]}(d)$, and a one-dimensional output layer.

The reference values for $u(T, 50, 50, \dots, 50)$ were calculated with the deep BSDE method of E, Han, and Jentzen [27]).

3.3. Allen–Cahn-type equations. Next, we approximate solutions of high-dimensional Allen–Cahn-type equations with a cubic nonlinearity of the form

$$(3.8) \quad \frac{\partial}{\partial t} u(t, x) = \Delta_x u(t, x) + u(t, x) - [u(t, x)]^3, \quad (t, x) \in (0, T] \times \mathbb{R}^d,$$

with initial condition $u(0, x) = \arctan(\max_{i \in \{1, 2, \dots, d\}} x_i)$; see also Beck, E, and Jentzen [3, section 4.1], E, Han, and Jentzen [27, section 4.2], E et al. [29, section 3.4], and Han, Jentzen, and E [47, section 3.3] for further numerical results for (3.8).

Table 3 lists deep splitting approximations of $u(T, 0, \dots, 0)$ for $T = 0.3$, $N = 10$, and different values of d . We chose $\mu(x) = (0, 0, \dots, 0) \in \mathbb{R}^d$, $\sigma(x) = \sqrt{2} \text{Id}_{\mathbb{R}^d \times d}$, and $f(x, y, z) = y - y^3$, $(x, y, z) \in \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d$. We used $M = 500$ and

$$(3.9) \quad \gamma_m = 10^{-1} \mathbb{1}_{[0,300]}(m) + 10^{-2} \mathbb{1}_{(300,400]}(m) + 10^{-3} \mathbb{1}_{(400,500]}(m).$$

We set $\xi^{n,m,j} = (0, 0, \dots, 0) \in \mathbb{R}^d$ for all $(n, m, j) \in \{1, 2, \dots, N\} \times \{0, 1, \dots, M-1\} \times \mathbb{N}$ and used feedforward neural networks with a d -dimensional input layer, two hidden layers of dimension $d + 10$, and a one-dimensional output layer.

The reference values for $u(T, 0, 0, \dots, 0)$ were calculated with the multilevel Picard method; see, e.g., [5, 28, 29, 56, 57, 58].

3.4. Semilinear heat equations. In this subsection, we calculate approximate solutions of semilinear heat equations of the form

$$(3.10) \quad \frac{\partial}{\partial t} u(t, x) = \Delta_x u(t, x) + \frac{1 - |u(t, x)|^2}{1 + |u(t, x)|^2}, \quad (t, x) \in (0, T] \times \mathbb{R}^d,$$

with initial condition $u(0, x) = 5/(10 + 2\|x\|_{\mathbb{R}^d}^2)$.

Table 4 shows deep splitting approximations of $u(T, 0, \dots, 0)$ for $T = 0.3$, $N = 20$, and different values of d . We derived them with $\mu(x) = (0, 0, \dots, 0) \in \mathbb{R}^d$, $\sigma(x) = \sqrt{2} \text{Id}_{\mathbb{R}^d \times d}$, and $f(x, y, z) = -\|z\|_{\mathbb{R}^d}^2$ for $(x, y, z) \in \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d$. We used $M = 500$ and

$$(3.11) \quad \gamma_m = 10^{-1} \mathbb{1}_{[0,300]}(m) + 10^{-2} \mathbb{1}_{(300,400]}(m) + 10^{-3} \mathbb{1}_{(400,500]}(m).$$

TABLE 4

Deep splitting approximations of solutions of the semilinear heat equation (3.10) for $T = 0.3$, $N = 20$, and different d .

d	Expectation	Std. dev.	Ref. value	Rel. L^1 -error	Std. dev. rel. error	Avg. runtime
10	0.47138	0.00035606	0.47006	0.00282	0.00075749	46.7s
50	0.34584	0.00018791	0.34425	0.00462	0.00054586	46.7s
100	0.31783	0.00008298	0.31674	0.00343	0.00026198	47.4s
200	0.30210	0.00002238	0.30091	0.00394	0.00007436	48.1s
300	0.29654	0.00001499	0.29534	0.00406	0.00005075	48.3s
500	0.29200	0.00000611	0.29095	0.00361	0.00002099	48.5s
1000	0.28852	0.00000267	0.28753	0.00344	0.00000930	54.1s
5,000	0.28569	0.00000042	0.28469	0.00352	0.00000148	286.3s
10,000	0.28533	0.00000048	0.28433	0.00353	0.00000170	1013.0s

TABLE 5

Deep splitting approximations of solutions of the semilinear heat equation (3.10) for $d = 100$, $T = 0.3$, and different N .

N	Expectation	Std. dev.	Ref. value	Rel. L^1 -error	Std. dev. rel. error	Avg. runtime
1	0.33820	0.00005021	0.31674	0.06777	0.00015852	0.5s
2	0.32836	0.00006189	0.31674	0.03669	0.00019539	2.2s
4	0.32259	0.00003834	0.31674	0.01848	0.00012104	5.7s
8	0.31969	0.00006929	0.31674	0.00931	0.00021878	12.8s
16	0.31810	0.00006968	0.31674	0.00430	0.00021998	27.1s
20	0.31783	0.00010565	0.31674	0.00345	0.00033354	34.3s
32	0.31739	0.00006564	0.31674	0.00206	0.00020723	55.6s

We set $\xi^{n,m,j} = (0, 0, \dots, 0) \in \mathbb{R}^d$ for all $(n, m, j) \in \{1, 2, \dots, N\} \times \{0, 1, \dots, M-1\} \times \mathbb{N}$ and used feedforward neural networks with a d -dimensional input layer, two hidden layers of dimension $d + 10$, and a one-dimensional output layer. The reference values were computed with the multilevel Picard method; see, e.g., [5, 28, 29, 56, 57, 58]).

To illustrate how the accuracy of the deep splitting method depends on the number of time steps N , we report in Table 5 deep splitting approximations of $u(T, 0, \dots, 0)$ for $d = 100$, $T = 0.3$, and different values of N . It can be seen that the decrease of the relative L^1 -error is approximately linear in N . For theoretical convergence results for the deep splitting method, we refer the reader to Beck et al. [4, Theorem 2].

3.5. Sine-Gordon-type equations. As a last example, we approximate solutions of sine-Gordon-type equations of the form

$$(3.12) \quad \frac{\partial}{\partial t} u(t, x) = \Delta_x u(t, x) + \sin(u(t, x)), \quad (t, x) \in (0, T] \times \mathbb{R}^d,$$

with initial condition $u(0, x) = 5/(10+2\|x\|_{\mathbb{R}^d}^2)$.

Table 6 shows deep splitting approximations of $u(T, 0, \dots, 0)$ for $T = 0.3$, $N = 20$, and different values of d . We chose $\mu(x) = (0, 0, \dots, 0) \in \mathbb{R}^d$, $\sigma(x) = \sqrt{2} \operatorname{Id}_{\mathbb{R}^d \times d}$, and $f(x, y, z) = \sin(y)$, $(x, y, z) \in \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d$. We used $M = 1000$ and

$$(3.13) \quad \gamma_m = 10^{-1} \mathbb{1}_{[0,250]}(m) + 10^{-2} \mathbb{1}_{(250,500]}(m) + 10^{-3} \mathbb{1}_{(500,750]}(m) + 10^{-4} \mathbb{1}_{(750,1000]}(m).$$

We set $\xi^{n,m,j} = (0, 0, \dots, 0) \in \mathbb{R}^d$ for every $(n, m, j) \in \{1, 2, \dots, N\} \times \{0, 1, \dots, M-1\} \times \mathbb{N}$ and used feedforward neural networks with a d -dimensional input layer, two hidden layers of dimension $d + 50$, and a one-dimensional output layer.

The reference values for $u(T, 0, 0, \dots, 0)$ were computed with the multilevel Picard method; see, e.g., [28, 29, 56, 57, 58].

TABLE 6

Deep splitting approximations of the solution of the sine-Gordon-type equation (3.12) for $T = 0.3$, $N = 20$, and different d .

d	Expectation	Std. dev.	Ref. value	Rel. L^1 -error	Std. dev. rel. error	Avg. runtime
10	0.3218822	0.00069331	0.3229470	0.0032972	0.0021468	55.6s
50	0.0990598	0.00013433	0.0993633	0.0030544	0.0013519	55.4s
100	0.0526955	0.00005390	0.0528368	0.0026741	0.0010202	55.6s
200	0.0271860	0.00001617	0.0272410	0.0020176	0.0005936	56.2s
300	0.0183162	0.00001010	0.0183617	0.0024765	0.0005502	55.8s
500	0.0110819	0.00000428	0.0111071	0.0022647	0.0003851	54.7s
1000	0.0055775	0.00000090	0.0055896	0.0021697	0.0001613	58.1s
5,000	0.0011209	0.00000012	0.0011231	0.0019464	0.0001070	332.5s
10,000	0.0005608	0.00000004	0.0005621	0.0022495	0.0000720	1083.4s

4. Comparison with other methods. In this section, we compare the deep splitting method with the deep BSDE method of E, Han, and Jentzen [27] and the multilevel Picard method of [5, 56], which also have been shown to produce good results in approximating solutions of high-dimensional PDEs.

Figure 1 shows estimated relative L^1 -approximation errors as a function of the number of one-dimensional standard normal random variables used by the three methods to approximate $u(T, 0, \dots, 0)$ for the solution $u: [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ of the sine-Gordon-type equation (3.12) for $d = 10$ and $T = 0.3$. It can be seen that for this particular example, the three methods yield comparable results. But it has to be noted that the deep splitting algorithm and the deep BSDE method both involve different hyperparameters, which, for good results, have to be fine-tuned depending on the form and the parameters of the PDE. Figure 1 just shows results for particular implementations of the three methods.

Generally, since the deep splitting method uses neural networks to approximate the solution $u: [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ of a PDE on a time grid $0 = t_0 < t_1 < \dots < t_N = T$, it can learn approximations of $u(t_n, x)$, $n \in \{1, 2, \dots, N\}$, simultaneously for all $x \in \mathbb{R}^d$. Similarly, the deep BSDE method can be implemented so that it approximates $u(T, x)$ directly for all $x \in \mathbb{R}^d$. However, for a temporal discretization with N subintervals, it then needs to train N neural networks at the same time, whereas the deep splitting method trains one network after the other. So even if the two approaches need a similar total number of one-dimensional standard normal random variables to achieve a given accuracy, the deep splitting method can handle larger problems since it decomposes them into smaller computational tasks which can be solved successively.

The advantage of the multilevel Picard method is that there exist theoretical bounds on the computational effort needed for a given approximation accuracy; see e.g., [28, 56, 57, 58]. However, the method needs to calculate approximations of $u(t, x)$ for different space-time points (t, x) separately and becomes impractical for large t .

5. Conclusion. In this paper, we have developed a new numerical method to approximate solutions of high-dimensional nonlinear parabolic PDEs. It splits the differential operator into a linear and a nonlinear part and uses deep learning together with the Feynman–Kac formula to iteratively solve linear approximations of the equation over small time intervals. This breaks the PDE approximation task into smaller problems that can be solved successively. As a consequence, the approach can be applied to extremely high dimensional nonlinear PDEs. We have tested the method on HJB equations, nonlinear Black–Scholes equations, Allen–Cahn-type equations, and semilinear heat equations as well as sine-Gordon-type equations. In all cases, it has

FIG. 1. Estimated relative L^1 -errors as a function of the number of one-dimensional standard normal random variables used by the deep splitting algorithm, the deep BSDE approach of E, Han, and Jentzen [27], and the multilevel Picard method of [5, 56] for the sine-Gordon-type equation (3.12) with $d = 10$ and $T = 0.3$.

produced accurate results in high dimensions with short run times.

Acknowledgment. We are grateful to Adam Andersson for fruitful discussions.

REFERENCES

- [1] V. BALLY AND G. PAGÈS, *A quantization algorithm for solving multi-dimensional discrete-time optimal stopping problems*, Bernoulli, 9 (2003), pp. 1003–1049.
- [2] C. BECK, S. BECKER, P. GROHS, N. JAAFARI, AND A. JENTZEN, *Solving the Kolmogorov PDE by means of deep learning*, J. Sci. Comput., 88 (2021), 73.
- [3] C. BECK, W. E, AND A. JENTZEN, *Machine learning approximation algorithms for high-dimensional fully nonlinear partial differential equations and second-order backward stochastic differential equations*, J. Nonlinear Sci., 29 (2019), pp. 1563–1619.
- [4] C. BECK, M. HUTZENTHALER, A. JENTZEN, AND B. KUCKUCK, *An Overview on Deep Learning-Based Approximation Methods for Partial Differential Equations*, preprint, <https://arxiv.org/abs/2012.12348>, 2020.
- [5] S. BECKER, R. BRAUNWARTH, M. HUTZENTHALER, A. JENTZEN, AND P. VON WURSTEMBERGER, *Numerical simulations for full history recursive multilevel Picard approximations for systems of high-dimensional partial differential equations*, Commun. Comput. Phys., 28 (2020), pp. 2109–2138.
- [6] S. BECKER, P. CHERIDITO, AND A. JENTZEN, *Deep optimal stopping*, J. Mach. Learn. Res., 20 (2019), 74.
- [7] C. BENDER AND R. DENK, *A forward scheme for backward SDEs*, Stochastic Process. Appl., 117 (2007), pp. 1793–1812.
- [8] C. BENDER, N. SCHWEIZER, AND J. ZHUO, *A primal-dual algorithm for BSDEs*, Math. Finance, 27 (2017), pp. 866–901.
- [9] Y. BENGIO, *Learning deep architectures for AI*, Found. Trends Mach. Learn., 2 (2009), pp. 1–127.
- [10] J. BERG AND K. NYSTRÖM, *A unified deep artificial neural network approach to partial differential equations in complex geometries*, Neurocomputing, 317 (2018), pp. 28–41.
- [11] J. BERNER, P. GROHS, AND A. JENTZEN, *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 J.

- Math. Data Sci., 2 (2020), pp. 631–657, <https://doi.org/10.1137/19M125649X>.
- [12] B. BOUCHARD AND N. TOUZI, *Discrete-time approximation and Monte-Carlo simulation of backward stochastic differential equations*, Stochastic Process. Appl., 111 (2004), pp. 175–206.
 - [13] D. BRAESS, *Finite Elements: Theory, Fast Solvers, and Applications in Elasticity Theory*, 3rd ed., Cambridge University Press, Cambridge, UK, 2007.
 - [14] Q. CHAN-WAI-NAM, J. MIKAEL, AND X. WARIN, *Machine learning for semi linear PDEs*, J. Sci. Comput., 79 (2019), pp. 1667–1712.
 - [15] J.-F. CHASSAGNEUX, *Linear multistep schemes for BSDEs*, SIAM J. Numer. Anal., 52 (2014), pp. 2815–2836, <https://doi.org/10.1137/120902951>.
 - [16] J.-F. CHASSAGNEUX AND D. CRISAN, *Runge-Kutta schemes for backward stochastic differential equations*, Ann. Appl. Probab., 24 (2014), pp. 679–720.
 - [17] J.-F. CHASSAGNEUX AND A. RICHOU, *Numerical stability analysis of the Euler scheme for BSDEs*, SIAM J. Numer. Anal., 53 (2015), pp. 1172–1193, <https://doi.org/10.1137/140977047>.
 - [18] J.-F. CHASSAGNEUX AND A. RICHOU, *Numerical simulation of quadratic BSDEs*, Ann. Appl. Probab., 26 (2016), pp. 262–304.
 - [19] D. CRISAN AND K. MANOLARAKIS, *Probabilistic methods for semilinear partial differential equations. Applications to finance*, M2AN Math. Model. Numer. Anal., 44 (2010), pp. 1107–1133.
 - [20] D. CRISAN AND K. MANOLARAKIS, *Solving backward stochastic differential equations using the cubature method: Application to nonlinear pricing*, SIAM J. Financial Math., 3 (2012), pp. 534–571, <https://doi.org/10.1137/090765766>.
 - [21] D. CRISAN AND K. MANOLARAKIS, *Second order discretization of backward SDEs and simulation with the cubature method*, Ann. Appl. Probab., 24 (2014), pp. 652–678.
 - [22] D. CRISAN, K. MANOLARAKIS, AND N. TOUZI, *On the Monte Carlo simulation of BSDEs: An improvement on the Malliavin weights*, Stochastic Process. Appl., 120 (2010), pp. 1133–1158.
 - [23] T. DECK AND S. KRUSE, *Parabolic differential equations with unbounded coefficients—a generalization of the parametrix method*, Acta Appl. Math., 74 (2002), pp. 71–91.
 - [24] F. DELARUE AND S. MENOZZI, *A forward-backward stochastic algorithm for quasi-linear PDEs*, Ann. Appl. Probab., 16 (2006), pp. 140–184.
 - [25] P. DÖRSEK, *Semigroup splitting and cubature approximations for the stochastic Navier–Stokes equations*, SIAM J. Numer. Anal., 50 (2012), pp. 729–746, <https://doi.org/10.1137/110833841>.
 - [26] J. DOUGLAS, JR., J. MA, AND P. PROTTER, *Numerical methods for forward-backward stochastic differential equations*, Ann. Appl. Probab., 6 (1996), pp. 940–968.
 - [27] 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.
 - [28] W. E, M. HUTZENTHALER, A. JENTZEN, AND T. KRUSE, *Multilevel Picard iterations for solving smooth semilinear parabolic heat equations*, Partial Differ. Equ. Appl., to appear.
 - [29] W. E, M. HUTZENTHALER, A. JENTZEN, AND T. KRUSE, *On multilevel Picard numerical approximations for high-dimensional nonlinear parabolic partial differential equations and high-dimensional nonlinear backward stochastic differential equations*, J. Sci. Comput., 79 (2019), pp. 1534–1571.
 - [30] W. E AND B. YU, *The deep Ritz method: A deep learning-based numerical algorithm for solving variational problems*, Commun. Math. Stat., 6 (2018), pp. 1–12.
 - [31] D. ELBRÄCHTER, P. GROHS, A. JENTZEN, AND C. SCHWAB, *DNN expression rate analysis of high-dimensional PDEs: Application to option pricing*, Constr. Approx., (2021), <https://doi.org/10.1007/s00365-021-09541-6>.
 - [32] A.-M. FARAHMAND, S. NABI, AND D. N. NIKOVSKI, *Deep reinforcement learning for partial differential equation control*, in Proceedings of the American Control Conference (ACC), 2017, pp. 3120–3127.
 - [33] M. FUJII, A. TAKAHASHI, AND M. TAKAHASHI, *Asymptotic expansion as prior knowledge in deep learning method for high dimensional BSDEs*, Asia-Pac., 26 (2019), pp. 391–408.
 - [34] X. GLOROT AND Y. BENGIO, *Understanding the difficulty of training deep feedforward neural networks*, in Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, 2010, pp. 249–256.
 - [35] E. GOBET AND C. LABART, *Solving BSDE with adaptive control variate*, SIAM J. Numer. Anal., 48 (2010), pp. 257–277, <https://doi.org/10.1137/090755060>.
 - [36] E. GOBET AND J.-P. LEMOR, *Numerical Simulation of BSDEs Using Empirical Regression*

- Methods: Theory and Practice*, preprint, <https://arxiv.org/abs/0806.4447>, 2008.
- [37] E. GOBET, J.-P. LEMOR, AND X. WARIN, *A regression-based Monte Carlo method to solve backward stochastic differential equations*, Ann. Appl. Prob., 15 (2005), pp. 2172–2202.
 - [38] E. GOBET, J. G. LÓPEZ-SALAS, P. TURKEDJIEV, AND C. VÁZQUEZ, *Stratified regression Monte-Carlo scheme for semilinear PDEs and BSDEs with large scale parallelization on GPUs*, SIAM J. Sci. Comput., 38 (2016), pp. C652–C677, <https://doi.org/10.1137/16M106371X>.
 - [39] E. GOBET AND P. TURKEDJIEV, *Approximation of backward stochastic differential equations using Malliavin weights and least-squares regression*, Bernoulli, 22 (2016), pp. 530–562.
 - [40] E. GOBET AND P. TURKEDJIEV, *Linear regression MDP scheme for discrete backward stochastic differential equations under general conditions*, Math. Comp., 85 (2016), pp. 1359–1391.
 - [41] L. GOUDENEGE, A. MOLENT, AND A. ZANETTE, *Machine learning for pricing American options in high-dimensional Markovian and non-Markovian models*, Quant. Finance, 20 (2020), pp. 573–591.
 - [42] W. GRECKSCH AND H. LISEI, *Approximation of stochastic nonlinear equations of Schrödinger type by the splitting method*, Stochastic Anal. Appl., 31 (2013), pp. 314–335.
 - [43] P. GROHS, F. HORNING, A. JENTZEN, AND P. VON WURSTENBERGER, *A proof that artificial neural networks overcome the curse of dimensionality in the numerical approximation of Black-Scholes partial differential equations*, Mem. Amer. Math. Soc., to appear.
 - [44] I. GYÖNGY AND N. KRYLOV, *On the rate of convergence of splitting-up approximations for SPDEs*, in Stochastic Inequalities and Applications, Birkhäuser, Basel, 2003, pp. 301–321.
 - [45] I. GYÖNGY AND N. KRYLOV, *On the splitting-up method and stochastic partial differential equations*, Ann. Probab., 31 (2003), pp. 564–591.
 - [46] M. HAIRER, M. HUTZENTHALER, AND A. JENTZEN, *Loss of regularity for Kolmogorov equations*, Ann. Probab., 43 (2015), pp. 468–527.
 - [47] 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.
 - [48] J. HAN AND J. LONG, *Convergence of the deep BSDE method for coupled FBSDEs*, Probab. Uncertain. Quant. Risk, 5 (2020), pp. 1–33.
 - [49] P. HENRY-LABORDÈRE, *Counterparty Risk Valuation: A Marked Branching Diffusion Approach*, preprint, <https://arxiv.org/abs/1203.2369>, 2012.
 - [50] P. HENRY-LABORDÈRE, *Deep Primal-Dual Algorithm for BSDEs: Applications of Machine Learning to CVA and IM*, preprint, SSRN-id3071506, 2017.
 - [51] P. HENRY-LABORDÈRE, N. OUDJANE, X. TAN, N. TOUZI, AND X. WARIN, *Branching diffusion representation of semilinear PDEs and Monte Carlo approximation*, Ann. Inst. Henri Poincaré Probab. Stat., 55 (2019), pp. 184–210.
 - [52] P. HENRY-LABORDÈRE, X. TAN, AND N. TOUZI, *A numerical algorithm for a class of BSDEs via the branching process*, Stochastic Process. Appl., 124 (2014), pp. 1112–1140.
 - [53] T. P. HUIJSKENS, M. J. RUIJTER, AND C. W. OOSTERLEE, *Efficient numerical Fourier methods for coupled forward-backward SDEs*, J. Comput. Appl. Math., 296 (2016), pp. 593–612.
 - [54] C. HURÉ, H. PHAM, AND X. WARIN, *Some Machine Learning Schemes for High-Dimensional Nonlinear PDEs*, preprint, <https://arxiv.org/abs/1902.01599v1>, 2019.
 - [55] M. HUTZENTHALER, A. JENTZEN, T. KRUSE, AND T. A. NGUYEN, *A proof that rectified deep neural networks overcome the curse of dimensionality in the numerical approximation of semilinear heat equations*, Partial Differ. Equ. Appl., 1 (2020), pp. 1–34.
 - [56] M. HUTZENTHALER, A. JENTZEN, T. KRUSE, T. A. NGUYEN, AND P. VON WURSTENBERGER, *Overcoming the curse of dimensionality in the numerical approximation of semilinear parabolic partial differential equations*, Proc. A, 476 (2020), 20190630.
 - [57] M. HUTZENTHALER, A. JENTZEN, AND P. VON WURSTENBERGER, *Overcoming the curse of dimensionality in the approximative pricing of financial derivatives with default risks*, Electron. J. Probab., 25 (2020), pp. 1–73.
 - [58] M. HUTZENTHALER AND T. KRUSE, *Multilevel Picard approximations of high-dimensional semilinear parabolic differential equations with gradient-dependent nonlinearities*, SIAM J. Numer. Anal., 58 (2020), pp. 929–961, <https://doi.org/10.1137/17M1157015>.
 - [59] S. IOFFE AND C. SZEGEDY, *Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift*, preprint, <https://arxiv.org/abs/1502.03167>, 2015.
 - [60] A. JACQUIER AND M. OUMGARI, *Deep PPDEs for Rough Local Stochastic Volatility*, preprint, <https://arxiv.org/abs/1906.02551v1>, 2019.
 - [61] A. JENTZEN, D. SALIMOVA, AND T. WELTI, *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*, Commun. Math. Sci., 19 (2021), pp. 1167–1205.
 - [62] D. KINGMA AND J. BA, *Adam: A Method for Stochastic Optimization*, preprint, <https://arxiv.org/abs/1412.0441>, 2014.

- org/abs/1412.6980, 2014.
- [63] A. KLENKE, *Probability Theory: A Comprehensive Course*, 2nd ed., Universitext, Springer, London, 2014.
 - [64] P. E. KLOEDEN AND E. PLATEN, *Numerical Solution of Stochastic Differential Equations*, Appl. Math. (N. Y.) 23, Springer-Verlag, Berlin, 1992.
 - [65] N. V. KRYLOV, *Lectures on Elliptic and Parabolic Equations in Hölder Spaces*, Grad. Stud. Math. 12, American Mathematical Society, Providence, RI, 1996.
 - [66] N. V. KRYLOV, *On Kolmogorov's equations for finite-dimensional diffusions*, in Stochastic PDE's and Kolmogorov Equations in Infinite Dimensions (Cetraro, 1998), Lecture Notes in Math. 1715, Springer, Berlin, 1999, pp. 1–63.
 - [67] G. KUTYNIK, P. PETERSEN, M. RASLAN, AND R. SCHNEIDER, *A Theoretical Analysis of Deep Neural Networks and Parametric PDEs*, <https://arxiv.org/abs/1904.00377>, 2019.
 - [68] C. LABART AND J. LELONG, *A parallel algorithm for solving BSDEs*, Monte Carlo Methods Appl., 19 (2013), pp. 11–39.
 - [69] S. LARSSON AND V. THOMÉE, *Partial Differential Equations with Numerical Methods*, Texts Appl. Math. 45, Springer-Verlag, Berlin, 2003.
 - [70] Y. LECUN, Y. BENGIO, AND G. HINTON, *Deep learning*, Nature, 521 (2015), pp. 436–444.
 - [71] J.-P. LEMOR, E. GOBET, AND X. WARIN, *Rate of convergence of an empirical regression method for solving generalized backward stochastic differential equations*, Bernoulli, 12 (2006), pp. 889–916.
 - [72] A. LIONNET, G. DOS REIS, AND L. SZPRUCH, *Time discretization of FBSDE with polynomial growth drivers and reaction-diffusion PDEs*, Ann. Appl. Probab., 25 (2015), pp. 2563–2625.
 - [73] Z. LONG, Y. LU, X. MA, AND B. DONG, *PDE-Net: Learning PDEs from data*, in Proceedings of the International Conference on Machine Learning, 2018, pp. 3208–3216.
 - [74] K. O. LYE, S. MISHRA, AND D. RAY, *Deep learning observables in computational fluid dynamics*, J. Comput. Phys., 410 (2020), 109339.
 - [75] J. MA, P. PROTTER, J. SAN MARTÍN, AND S. TORRES, *Numerical method for backward stochastic differential equations*, Ann. Appl. Probab., 12 (2002), pp. 302–316.
 - [76] 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.
 - [77] J. MA AND J. YONG, *Forward-Backward Stochastic Differential Equations and Their Applications*, Lecture Notes in Math. 1702, Springer-Verlag, Berlin, 1999.
 - [78] M. MAGILL, F. QURESHI, AND H. DE HAAN, *Neural networks trained to solve differential equations learn general representations*, in Advances in Neural Information Processing Systems, MIT Press, Cambridge, MA, 2018, pp. 4075–4085.
 - [79] G. MARUYAMA, *Continuous Markov processes and stochastic equations*, Rend. Circ. Mat. Palermo (2), 4 (1955), pp. 48–90.
 - [80] H. P. MCKEAN, *Application of Brownian motion to the equation of Kolmogorov-Petrovskii-Piskunov*, Comm. Pure Appl. Math., 28 (1975), pp. 323–331.
 - [81] G. N. MIL'STEIN, *Approximate integration of stochastic differential equations*, Theory Probab. Appl., 19 (1975), pp. 557–562, <https://doi.org/10.1137/1119062>.
 - [82] G. N. MILSTEIN AND M. V. TRETYAKOV, *Numerical algorithms for forward-backward stochastic differential equations*, SIAM J. Sci. Comput., 28 (2006), pp. 561–582, <https://doi.org/10.1137/040614426>.
 - [83] G. N. MILSTEIN AND M. V. TRETYAKOV, *Discretization of forward-backward stochastic differential equations and related quasi-linear parabolic equations*, IMA J. Numer. Anal., 27 (2007), pp. 24–44.
 - [84] G. N. MILSTEIN AND M. V. TRETYAKOV, *Solving parabolic stochastic partial differential equations via averaging over characteristics*, Math. Comp., 78 (2009), pp. 2075–2106.
 - [85] H. PHAM, *Feynman-Kac representation of fully nonlinear PDEs and applications*, Acta Math. Vietnam., 40 (2015), pp. 255–269.
 - [86] M. RAISSI, *Deep hidden physics models: Deep learning of nonlinear partial differential equations*, J. Mach. Learn. Res., 19 (2018), pp. 932–955.
 - [87] M. J. RUIJTER AND C. W. OOSTERLEE, *A Fourier cosine method for an efficient computation of solutions to BSDEs*, SIAM J. Sci. Comput., 37 (2015), pp. A859–A889, <https://doi.org/10.1137/130913183>.
 - [88] M. J. RUIJTER AND C. W. OOSTERLEE, *Numerical Fourier method and second-order Taylor scheme for backward SDEs in finance*, Appl. Numer. Math., 103 (2016), pp. 1–26.
 - [89] A. RUSZCZYŃSKI AND J. YAO, *A Dual Method for Backward Stochastic Differential Equations with Application to Risk Valuation*, preprint, <https://arxiv.org/abs/1701.06234>, 2017.
 - [90] J. SIRIGNANO AND K. SPILIOPOULOS, *DGM: A deep learning algorithm for solving partial dif-*

- ferential equations*, J. Comput. Phys., 375 (2018), pp. 1339–1364.
- [91] V. V. SKOROKHOD, *Branching diffusion processes*, Theory Probab. Appl., 9 (1964), pp. 445–449, <https://doi.org/10.1137/1109059>.
 - [92] D. W. STROOCK, *Lectures on Topics in Stochastic Differential Equations*, Tata Inst. Fund. Res. Lectures on Math. and Phys. 68, Tata Institute of Fundamental Research, Bombay; Springer-Verlag, Berlin, New York, 1982.
 - [93] V. THOMÉE, *Galerkin Finite Element Methods for Parabolic Problems*, Springer Ser. Comput. Math. 25, Springer-Verlag, Berlin, 1997.
 - [94] P. TURKEDJIEV, *Two algorithms for the discrete time approximation of Markovian backward stochastic differential equations under local conditions*, Electron. J. Probab., 20 (2015), 50.
 - [95] S. WATANABE, *On the branching process for Brownian particles with an absorbing boundary*, J. Math. Kyoto Univ., 4 (1965), pp. 385–398.
 - [96] J. ZHANG, *A numerical scheme for BSDEs*, Ann. Appl. Probab., 14 (2004), pp. 459–488.