

---

# Numerically Solving Parametric Families of High-Dimensional Kolmogorov Partial Differential Equations via Deep Learning

---

**Julius Berner\***

Faculty of Mathematics, University of Vienna  
Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria  
julius.berner@univie.ac.at

**Markus Dablander\***

Mathematical Institute, University of Oxford  
Andrew Wiles Building, OX2 6GG, Oxford, United Kingdom  
markus.dablander@maths.ox.ac.uk

**Philipp Grohs**

Faculty of Mathematics and Research Platform DataScience@UniVienna, University of Vienna  
Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria  
RICAM, Austrian Academy of Sciences  
Altenberger Straße 69, 4040 Linz, Austria  
philipp.grohs@univie.ac.at

## Abstract

We present a deep learning algorithm for the numerical solution of parametric families of high-dimensional linear Kolmogorov partial differential equations (PDEs). Our method is based on reformulating the numerical approximation of a whole family of Kolmogorov PDEs as a single statistical learning problem using the Feynman-Kac formula. Successful numerical experiments are presented, which empirically confirm the functionality and efficiency of our proposed algorithm in the case of heat equations and Black-Scholes option pricing models parametrized by affine-linear coefficient functions. We show that a single deep neural network trained on simulated data is capable of learning the solution functions of an entire family of PDEs on a full space-time region. Most notably, our numerical observations and theoretical results also demonstrate that the proposed method does not suffer from the curse of dimensionality, distinguishing it from almost all standard numerical methods for PDEs.

## 1 Introduction

Linear parabolic partial differential equations (PDEs) of the form

$$\frac{\partial u_\gamma}{\partial t} = \frac{1}{2} \text{Trace}(\sigma_\gamma [\sigma_\gamma]^* \nabla_x^2 u_\gamma) + \langle \mu_\gamma, \nabla_x u_\gamma \rangle, \quad u_\gamma(x, 0) = \varphi_\gamma(x), \quad (1)$$

are referred to as Kolmogorov PDEs, see [23] for a thorough study of their mathematical properties. Throughout this paper, the functions

$$\varphi_\gamma : \mathbb{R}^d \rightarrow \mathbb{R} \quad (\text{initial condition}) \quad \text{and} \quad \sigma_\gamma : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}, \quad \mu_\gamma : \mathbb{R}^d \rightarrow \mathbb{R}^d \quad (\text{coefficient maps})$$

---

\*Equal contribution.

are continuous, and are implicitly determined by a real parameter vector  $\gamma \in D$ , whereby  $D$  is a compact set in Euclidean space.

Kolmogorov PDEs frequently appear in applications in a broad variety of scientific disciplines, including physics and financial engineering [9, 42, 58]. In particular, note that the heat equation from physical modelling as well as the widely-known Black-Scholes equation from computational finance are important special cases of Equation (1). Typically, one is interested in finding the (viscosity) solution<sup>2</sup>

$$u_\gamma : [v, w]^d \times [0, T] \rightarrow \mathbb{R}$$

of Equation (1) on a predefined space-time region of the form  $[v, w]^d \times [0, T]$ . In almost all cases, however, Kolmogorov PDEs cannot be solved explicitly. Furthermore, standard numerical solution algorithms for PDEs, in particular those based on a discretization of the considered domain, are known to suffer from the so-called *curse of dimensionality*<sup>3</sup>, meaning that their computational cost grows exponentially in the dimension of the domain [2, 51]. The development of new, computationally efficient methods for the numerical solution of Kolmogorov PDEs is therefore of high interest for applied scientists.

In this work, we present a novel deep learning algorithm capable of numerically approximating the solutions  $(u_\gamma)_{\gamma \in D}$  of a whole family of  $\gamma$ -parametrized Kolmogorov PDEs on a full space-time region. Specifically, our proposed method allows to train a single deep neural network

$$\Phi : D \times [v, w]^d \times [0, T] \rightarrow \mathbb{R} \quad (2)$$

to approximate the *parametric solution map*

$$\bar{u} : D \times [v, w]^d \times [0, T] \rightarrow \mathbb{R}, \quad (\gamma, x, t) \mapsto \bar{u}(\gamma, x, t) := u_\gamma(x, t), \quad (3)$$

of a family of  $\gamma$ -parametrized Kolmogorov PDEs on the generalized domain  $D \times [v, w]^d \times [0, T]$ . Most notably, we also theoretically investigate the associated approximation and generalization errors and rigorously show that our algorithm does not suffer from the curse of dimensionality with respect to the neural network size as well as the sample size. We emphasize that our empirical results strongly suggest that also the empirical risk minimization (ERM) algorithm, usually a variant of stochastic gradient descent, does not suffer from the curse of dimensionality but proving this is out of scope of this paper.

## 1.1 PDEs and Deep Learning: Current Research and Related Work

Interest in deep-learning based techniques for the numerical solution of PDEs has been growing rapidly in recent years [6, 24, 32, 47, 53, 56, 57]. This sharp rise in interest can partly be explained by the remarkable ability of deep neural networks to avoid incurring the curse of dimensionality when used to approximate PDE solutions in particular settings. More specifically, in some situations it has been possible to find theoretical upper bounds for the size of the required neural network architectures which do not depend exponentially on the dimension of the PDE [14, 20, 29, 50, 48]. This represents a rare and crucial advantage over classical finite difference and finite element methods, all of which typically cannot be used in high dimensions due to the resulting exponential explosion of required computational costs.

As a result of these successes, deep learning has recently been studied as a numerical solution technique for the more general group of parametric PDEs and their associated parametric solution maps [12, 27, 33, 36, 37, 50]. The investigation of the capabilities of deep artificial neural networks to learn parametric solution maps of specific parametrizable families of PDEs has become a new and active area of research. In this work, we provide novel theoretical and empirical results which, for the first time, demonstrate the viability of deep learning algorithms for the scalable solution of large classes of parametric Kolmogorov PDEs.

The formulation of the learning problem underlying our method is inspired by the work of Beck et al. [5]. There it is shown how deep neural networks can be used to numerically solve a non-parametric version of Equation (1) with fixed initial condition  $\varphi_\gamma$  and fixed coefficients maps  $\sigma_\gamma, \mu_\gamma$

<sup>2</sup>Viscosity solutions are the appropriate solution concept for a wide range of PDEs [10, 23]. Viscosity solutions are continuous, but not necessarily differentiable.

<sup>3</sup>The classical way to circumvent the curse of dimensionality has been the employment of stochastic Monte Carlo based methods, see e.g. [19]; these techniques, however, are only suitable to approximately compute the solution  $u_\gamma(x, t)$  at a single *fixed* space-time point  $(x, t) \in [v, w]^d \times [0, T]$ , limiting their usefulness in practice.

on a predefined space region  $[v, w]^d$  and at a predefined time slice  $T > 0$ . In other words, their non-parametric method allows to approximate the function

$$u_\gamma(\cdot, T) : [v, w]^d \rightarrow \mathbb{R}, \quad x \mapsto u_\gamma(x, T),$$

for fixed  $\gamma \in D$  by training a deep neural network with suitable simulated data of the form

$$(X, \varphi_\gamma(S_{\gamma, X, T})) \in [v, w]^d \times \mathbb{R}.$$

Here,  $X$  is uniformly drawn from the spatial hypercube  $[v, w]^d$  and the random vector  $S_{\gamma, X, T}$  is the value of the solution process  $(S_{\gamma, X, t})_{t \geq 0}$  of the stochastic differential equation (SDE)

$$dS_{\gamma, X, t} = \mu_\gamma(S_{\gamma, X, t})dt + \sigma_\gamma(S_{\gamma, X, t})dB_t, \quad S_{\gamma, X, 0} = X,$$

at time  $t = T$ , whereby  $(B_t)_{t \geq 0}$  is a standard  $d$ -dimensional Brownian motion.

The choice of training data is based on the following important identity, which under suitable regularity assumptions holds for all  $x \in [v, w]^d$ ,  $t \in [0, T]$ , and  $\gamma \in D$ :

$$u_\gamma(x, t) = \mathbb{E}[\varphi_\gamma(S_{\gamma, x, t})]. \quad (4)$$

Equality (4) is a version of the well-known *Feynman-Kac formula* which establishes a link between the theory of parabolic PDEs and the theory of stochastic differential equations [23]. Using the Feynman-Kac formula, one can show within the mathematical framework of empirical risk minimization [11, 55] that  $u_\gamma(\cdot, T)$  is in fact the solution of the supervised statistical learning problem defined by the predictor variable  $X$ , the target variable  $\varphi_\gamma(S_{\gamma, X, T})$ , and a standard quadratic loss function [5].

## 1.2 Novel Contribution

In this work, we introduce the first algorithm for the numerical solution of *parametric* Kolmogorov PDEs on a *whole* space-time region. No previous technique has achieved this degree of generality; all former methods for parametric Kolmogorov PDEs were either only capable of computing local solutions at single space-time points of the domain using Monte Carlo based approaches or did not employ deep neural networks and were thus not able to break the curse of dimensionality. Our technique is made possible by constructing a suitable supervised learning problem via a nontrivial application of the Feynman-Kac formula (4), which involves random initial conditions and SDEs with random coefficients and stopping times. This reformulation of a broad class of parametric PDEs as learning problems provides a new theoretical framework to analyze the convergence behavior of deep learning algorithms. Building upon this framework, we prove theoretical guarantees for the computational performance of our technique and, to the best of our knowledge, establish the first combined approximation and generalization results for parametric PDEs.

Note that the parametric nature of the presented algorithm opens up the novel possibility to study changes in the potentially high-dimensional solution manifold of Equation (1) as its initial conditions and coefficient maps vary with  $\gamma \in D$ . The study of such changes is commonly referred to as *sensitivity analysis*. Recall that the proposed method delivers a neural network  $\Phi$  which approximates the parametric PDE solution map, i.e.  $\Phi \approx \bar{u}$ . The partial derivatives of  $\Phi$  with respect to the parameter  $\gamma$ , the spatial variable  $x$ , and the time variable  $t$  can then be readily computed via automatic differentiation. Thus, the partial derivatives of  $\Phi$  become computationally accessible approximations of the partial derivatives of  $\bar{u}$ . The partial derivatives of  $\bar{u}$  in turn play an important role in a variety of widespread applications, such as in the computation of the ‘‘Greeks’’ associated with the Black-Scholes model from financial engineering (see Section 3.1).

Another highly relevant application area opened up by our method is the calibration of the usually unknown PDE coefficients  $\sigma_\gamma, \mu_\gamma$  using real-world data. After solving a parametric PDE with our technique, one can fit  $\gamma$  such that the PDE solution manifold best describes a real data set and additionally apply uncertainty quantification techniques if  $\gamma$  is modelled as a random variable.

Finally, we establish a new architecture and compare different learning schemes to provide future researchers with a robust framework for parametric PDEs, which are inherently less stable than their simpler non-parametric counterparts. Further, this work is complemented by an extendable implementation with the possibility of distributed training and hyperparameter optimization for the special use-cases of other researchers.

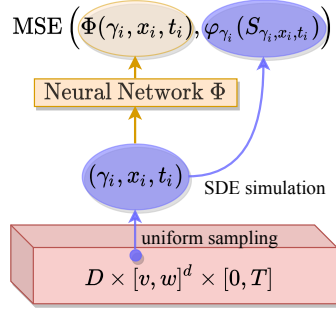


Figure 1: Illustration of the proposed supervised learning problem with predictor variable  $\Lambda$  and target variable  $\varphi_\Gamma(S_{\Gamma, X, \mathcal{T}})$ .

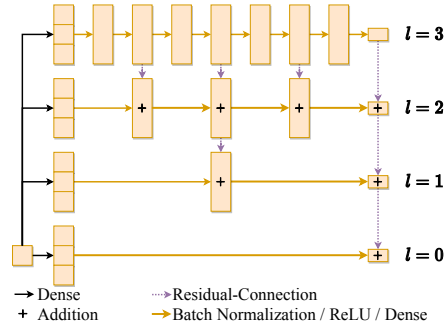


Figure 2: Illustration of the Multilevel architecture for  $L = 4, q = 3$ .

## 2 Algorithm

The key idea of the presented algorithm is to describe the parametric solution map (3) of the  $\gamma$ -parametrized Kolmogorov PDE (1) as the regression function of an appropriately chosen supervised statistical learning problem. One can then use simulated training data in order to learn  $\bar{u}$  by means of deep learning. Inspired by the Feynman-Kac formula (4), we construct a new statistical learning problem via a uniformly distributed predictor variable and a statistically dependent target variable:

$$\Lambda := (\Gamma, X, \mathcal{T}) \in D \times [v, w]^d \times [0, T] \quad (\text{predictor}) \quad \text{and} \quad Y := \varphi_\Gamma(S_\Lambda) \in \mathbb{R} \quad (\text{target}).$$

The random vector  $S_\Lambda$  is defined as the value of the solution process  $(S_{\Gamma, X, t})_{t \geq 0}$  of the  $\Gamma$ -parametrized stochastic differential equation

$$dS_{\Gamma, X, t} = \mu_\Gamma(S_{\Gamma, X, t})dt + \sigma_\Gamma(S_{\Gamma, X, t})dB_t, \quad S_{\Gamma, X, 0} = X, \quad (5)$$

at the (random) stopping time  $t = \mathcal{T}$ . For suitable regularity assumptions, the Feynman-Kac formula (4) then ensures that

$$\mathbb{E}[Y \mid \Lambda = (\gamma, x, t)] = \mathbb{E}[\varphi_\Gamma(S_\Lambda) \mid \Lambda = (\gamma, x, t)] = \mathbb{E}[\varphi_\gamma(S_{\gamma, x, t})] = u_\gamma(x, t) = \bar{u}(\gamma, x, t).$$

This shows that the minimizer of the corresponding statistical learning problem with quadratic loss function is indeed the parametric Kolmogorov PDE solution map, see Section A.1 in the appendix for the precise assumptions and a detailed proof.

**Theorem 1** (Learning Problem). *It holds that the parametric solution map  $\bar{u}$  is the unique minimizer of the statistical learning problem*

$$\min_f \mathbb{E}[(f(\Lambda) - Y)^2]. \quad (6)$$

Restricting to a hypothesis space of suitable neural networks  $\mathcal{H}$  and minimizing the empirical mean squared error (MSE) loss corresponding to (6), we arrive at the feasible supervised ERM problem

$$\min_{\Phi \in \mathcal{H}} \frac{1}{s} \sum_{i=1}^s (\Phi(\lambda_i) - y_i)^2 \quad (7)$$

where  $((\lambda_i, y_i))_{i=1}^s$  are realizations of i.i.d. samples drawn from the distribution of  $(\Lambda, Y)$ . Typically, this problem is then solved by a variant of stochastic gradient descent [49]. The algorithm is graphically illustrated in Figure 1.

It is trivial to simulate i.i.d. samples of the predictor variable  $\Lambda$ , due to its uniform distribution. On the other hand i.i.d. samples of the target variable  $Y = \varphi_\Gamma(S_\Lambda)$  can be obtained via standard numerical SDE solution techniques without curse of dimensionality [35]. An example for such a technique is given by the *Euler-Maruyama approximation* with  $M \in \mathbb{N}$  equidistant steps  $(S_\Lambda^{M, m})_{m=0}^M$  which is defined by the following scheme:

$$S_\Lambda^{M, 0} = X \quad \text{and} \quad S_\Lambda^{M, m+1} = S_\Lambda^{M, m} + \mu_\Gamma(S_\Lambda^{M, m}) \frac{\mathcal{T}}{M} + \sigma_\Gamma(S_\Lambda^{M, m}) (B_{\frac{(m+1)\mathcal{T}}{M}} - B_{\frac{m\mathcal{T}}{M}}). \quad (8)$$

The following theorem shows that solving the learning problem with data simulated by the Euler-Maruyama scheme does indeed result in the expected approximation of the parametric PDE solution map  $\bar{u}$ , see Section A.1 in the appendix for a proof.

**Theorem 2** (Approximated Learning Problem). *The unique minimizer  $\bar{u}^M$  of the approximated statistical learning problem*

$$\min_f \mathbb{E}[(f(\Lambda) - Y^M)^2]$$

where  $Y^M := \varphi_\Gamma(S_\Lambda^{M,M})$  is simulated using the Euler-Maruyama scheme (8) with  $M \sim 1/\varepsilon^2$  equidistant steps satisfies that

$$\|\bar{u}^M - \bar{u}\|_{\mathcal{L}^\infty(D \times [v,w]^d \times [0,T])} \leq \varepsilon.$$

In other words, the approximation of the SDE solution  $S_\Lambda^{M,M} \approx S_\Lambda$  carries over to the approximation of the corresponding minimizer  $\bar{u}^M \approx \bar{u}$ . Therefore there exists no constraint of having to solve the SDE in (5) analytically. The ability to easily simulate artificial training data opens up the highly desirable capability to supply the learning algorithm with a potentially infinite stream of i.i.d. data samples. Instead of having to use a train/val/test split on a given finite data set, one can thus constantly simulate new data points on demand during training. Since the number of samples then grows at will parallel to the training process, the first epoch never finishes and every new gradient computation can be done on the basis of previously unseen data.

## 2.1 Example: Affine-Linear Coefficient Functions

In Section 3 we present numerical experiments based on our algorithm for the important special case where  $\sigma_\gamma$  and  $\mu_\gamma$  are affine-linear functions. Thus from now on let us assume that<sup>4</sup>

$$\begin{aligned} \sigma_\gamma : \mathbb{R}^d &\rightarrow \mathbb{R}^{d \times d}, & \sigma_\gamma(x) &= [\gamma_{\sigma,1x} | \dots | \gamma_{\sigma,dx}] + \gamma_{\sigma,d+1}, \\ \mu_\gamma : \mathbb{R}^d &\rightarrow \mathbb{R}^d, & \mu_\gamma(x) &= \gamma_{\mu,1x} + \gamma_{\mu,2}, \end{aligned}$$

are affine-linear functions, which are determined by parameter tuples of matrices and vectors

$$\gamma_\sigma \in D_\sigma \subseteq (\mathbb{R}^{d \times d})^{d+1}, \quad \gamma_\mu \in D_\mu \subseteq \mathbb{R}^{d \times d} \times \mathbb{R}^d.$$

The parameter sets  $D_\sigma$  and  $D_\mu$  are chosen to be compact. Together with a suitable compact parameter set  $D_\varphi \subseteq \mathbb{R}^k$  for the initial function  $\varphi_\gamma$ , we obtain

$$\gamma := (\gamma_\sigma, \gamma_\mu, \gamma_\varphi) \in D_\sigma \times D_\mu \times D_\varphi := D \subseteq (\mathbb{R}^{d \times d})^{d+1} \times (\mathbb{R}^{d \times d} \times \mathbb{R}^d) \times \mathbb{R}^k.$$

This leads to an input dimension of our neural network  $\Phi$  of

$$\dim_{\text{in}}(\Phi) = \dim(D \times [v,w]^d \times [0,T]) = d^3 + 2d^2 + 2d + 1 + k.$$

Kolmogorov PDEs with affine-linear coefficient functions regularly appear in applications; the heat equation from physics and the classical and generalized Black-Scholes equations from computational finance are important examples of Kolmogorov PDEs with affine-linear coefficient maps [13, 45]. Note that while affine-linear coefficient functions are important in practise, computationally fast to evaluate, and easy to parametrize, the presented method is not restricted to the case of affine-linear coefficients and can as well be used in a substantially more general setting.

## 3 Numerical Results

We implemented the framework described in Section 2 in PyTorch [43] and computed our results on a Nvidia DGX-1 using Tune [39] for experiment execution and hyperparameter optimization. In this section, we describe our setting and present four encouraging demonstrations of the performance of our algorithm.<sup>5</sup>

For the neural network  $\Phi$  we propose a *Multilevel architecture* which is inspired by multilevel techniques such as Multilevel Monte Carlo methods [18], network architectures in [26, 61], and the architecture for the squaring function used in the proofs of our theoretical results in Section A.1, see also [60, Figure 2c]. One can view the output of the network  $\Phi$  as a sum  $\sum_{l=0}^{L-1} \Phi_l$  of sub-networks

<sup>4</sup>We denote by  $[a_1 | \dots | a_d] \in \mathbb{R}^{d \times d}$  the horizontal concatenation of the vectors  $a_1, \dots, a_d \in \mathbb{R}^d$ .

<sup>5</sup>For the implementation details we refer the reader to Section A.2 in the appendix and the repository associated with this work on [https://github.com/juliusberner/deep\\_kolmogorov](https://github.com/juliusberner/deep_kolmogorov).

$\Phi_l$  with  $2^l$  layer each. We think of the shallow network  $\Phi_0$  as computing a coarse approximation of  $\bar{u}$  and of the deep networks  $\Phi_l$  (with  $l \geq 1$ ) as approximately learning the residuals  $\bar{u} - \sum_{i=0}^{l-1} \Phi_i$ . To facilitate optimization we normalize our inputs<sup>6</sup> and to enhance expressivity we add an initial layer which increases the width by a given factor  $q$ . The Multilevel architecture is depicted in Figure 2 and described in detail in Definition 1 in the appendix.

Our optimized hyperparameters as well as an ablation study of our architecture and training scheme can be found in Sections A.2 and A.3 in the appendix. For all our experiments we were able to stick to a similar setup which depicts its robustness and general applicability. This is also mirrored by the small standard deviations of our considered errors across independent runs in Tables 1, 2, 3, and 4. These tables report average runtimes (in seconds), average approximation errors, and their standard deviations w.r.t. 4 independent runs each 4000 gradient descent steps. As an evaluation metric we approximately computed  $\mathcal{L}^1$ -errors via Monte Carlo sampling, that is

$$\left\| \frac{\Phi(\Lambda) - \bar{u}(\Lambda)}{1 + |\bar{u}(\Lambda)|} \right\|_{\mathcal{L}^1} := \mathbb{E} \left[ \frac{|\Phi(\Lambda) - \bar{u}(\Lambda)|}{1 + |\bar{u}(\Lambda)|} \right] \approx \frac{1}{n} \sum_{i=1}^n \left( \frac{|\Phi(\lambda_i) - \bar{u}(\lambda_i)|}{1 + |\bar{u}(\lambda_i)|} \right) \quad (9)$$

with  $n \in \mathbb{N}$  realizations  $(\lambda_i)_{i=1}^n$  of i.i.d. samples drawn from the distribution of  $\Lambda$  (drawn independently of the training data in (7) and drawn independently for each evaluation step). This means that we always evaluate our model w.r.t. to the parametric solution map  $\bar{u}$  on unseen input data; if no closed-form solution for  $\bar{u}$  is available, as in the case of the Basket option in 3.2 below, we approximate  $\bar{u}(\lambda_i)$  pointwise via Monte Carlo sampling, i.e.

$$\bar{u}(\lambda_i) = \bar{u}(\gamma_i, x_i, t_i) = \mathbb{E}[\varphi_{\gamma_i}(S_{\gamma_i, x_i, t_i})] \approx \frac{1}{m} \sum_{j=1}^m \varphi_{\gamma_i}(s_j) \quad (10)$$

where  $(s_j)_{j=1}^m$  are realizations of i.i.d. samples drawn from the distribution of the Euler-Maruyama approximation  $S_{\lambda_i}^{M, M}$  (drawn independently of the training data in (7) and drawn independently for each point and each evaluation step). Note that (9) is invariant under scaling of the hypercubes and locally corresponds to relative errors where the solution  $\bar{u}$  is large and absolute errors where it is small, which in particular prevents division by zero.

### 3.1 Black-Scholes Options Pricing Model

Our first example shows that neural networks are capable of learning a parametric version of the highly-celebrated Black-Scholes option pricing model [9]. We consider a European put option which gives its owner the right, but not the obligation, to sell a single underlying financial asset at a specified strike price and at a given time. Formally, this corresponds to  $d = 1$  and

$$\sigma_\gamma(x) = \gamma_\sigma x, \quad \mu_\gamma(x) = 0, \quad \varphi_\gamma(x) = \max\{\gamma_\varphi - x, 0\}, \quad x \in \mathbb{R},$$

with<sup>7</sup>  $\gamma_\sigma \in D_\sigma \subseteq \mathbb{R}$  and  $\gamma_\varphi \in D_\varphi \subseteq \mathbb{R}$ . Effectively, this leads to an input dimension of our neural network  $\Phi$  of  $\dim_{\text{in}}(\Phi) = 4$ . In case of the present Black-Scholes model, the associated SDE in (5) can actually be solved explicitly; it gives rise to geometric Brownian motion with uniformly distributed volatility  $\Gamma_\sigma \in D_\sigma$  and stopping time  $\mathcal{T} \in [0, T]$ , i.e.

$$S_\Lambda = X e^{-0.5\mathcal{T}\Gamma_\sigma^2 + \sqrt{\mathcal{T}}\Gamma_\sigma N}, \quad N \sim \mathcal{N}(0, 1).$$

We thus obtain an analytic expression for the parametric PDE solution,

$$\bar{u}(\gamma, x, t) = \gamma_\varphi \Psi(h_\gamma(x, t) + \sqrt{t}\gamma_\sigma) - x \Psi(h_\gamma(x, t)),$$

and the partial derivatives, e.g.

$$\frac{\partial \bar{u}}{\partial \gamma_\sigma}(\gamma, x, t) = x \sqrt{t} \Psi'(-h_\gamma(x, t)),$$

where

$$\Psi(z) := \frac{1}{2} \left[ 1 + \operatorname{erf} \left( \frac{z}{\sqrt{2}} \right) \right] \quad \text{and} \quad h_\gamma(x, t) := -\frac{1}{\sqrt{t}\gamma_\sigma} \left[ \ln \left( \frac{x}{\gamma_\varphi} \right) + \frac{t\gamma_\sigma^2}{2} \right],$$

see [4, Section 13.7]. This analytic expression can be used to evaluate the performance of our algorithm. We point out that the partial derivatives of  $\bar{u}$  are crucial in option pricing and each of

<sup>6</sup>We know the underlying (uniform) distribution and therefore can normalize each input in an exact manner.

<sup>7</sup>Note that  $\sigma_\gamma = \sigma_{\gamma,1}$  in the formal framework described in Section 2.1 but here and in the following we use the natural identifications, e.g.  $D_\sigma \cong D_\sigma \times \{0\}$ .




Figure 3: Shows the average prediction error  $\frac{|\Phi(\gamma, \cdot, t) - \bar{u}(\gamma, \cdot, t)|}{1 + |\bar{u}(\gamma, \cdot, t)|}$  and its standard deviation at  $t = 0.5$ ,  $\gamma_\sigma = 0.35$ , and  $\gamma_\varphi = 11$ .




Figure 4: Shows the average error of the Vega  $\frac{|\frac{\partial \Phi}{\partial \gamma_\sigma}(\gamma, \cdot, t) - \frac{\partial \bar{u}}{\partial \gamma_\sigma}(\gamma, \cdot, t)|}{1 + |\frac{\partial \bar{u}}{\partial \gamma_\sigma}(\gamma, \cdot, t)|}$  and its standard deviation at  $t = 0.5$ ,  $\gamma_\sigma = 0.35$ , and  $\gamma_\varphi = 11$ .

them is associated with a distinct economic interpretation. They are often referred to as *Greeks* and describe the sensitivity of the option price w.r.t. different model parameters, see for instance [4, 52]. The most prominent Greeks are given by

$$\Delta = \frac{\partial \bar{u}}{\partial x}, \quad \text{Vega} = \frac{\partial \bar{u}}{\partial \gamma_\sigma}, \quad \Theta = -\frac{\partial \bar{u}}{\partial t}.$$

On the basis of the proposed algorithm, our neural network  $\Phi$  is capable of learning the parametric solution map  $\bar{u}$  of the above problem in 24000 gradient updates up to an average  $\mathcal{L}^1$ -error of 0.0011, see Table 1 and Figure 3. As expected, the partial derivatives of our network  $\Phi$  (computed via automatic differentiation) approximate the partial derivatives of  $\bar{u}$  as can be seen in Figure 4. Further evidence can be found in Figures 5, 6, 7, and 8 in the appendix. Even though the parametric PDE problem can be solved explicitly in this special case, we use this relatively simple example for the purpose of illustrating our algorithm in an intuitive setting. As we will see below, the proposed algorithm is by no means restricted to such basic examples and can be applied successfully to much more complex and high-dimensional problems as well.

### 3.2 Basket Put Option

In the following we show that we can obtain comparable results to Section 3.1 in the case of a considerably more complicated Basket put option pricing problem, where analytical solutions of the PDE and the SDE are lacking. In such cases, our algorithm allows practitioners to nevertheless gain valuable insights into the behaviour of the PDE solution manifold as input parameters vary. By means of our trained model  $\Phi$  one can easily compute sensitivity values  $\frac{\partial \Phi}{\partial \gamma} \approx \frac{\partial \bar{u}}{\partial \gamma}$ ,  $\frac{\partial \Phi}{\partial t} \approx \frac{\partial \bar{u}}{\partial t}$ , and  $\frac{\partial \Phi}{\partial x} \approx \frac{\partial \bar{u}}{\partial x}$  via automatic differentiation or fit the parameter  $\gamma$  to a real-world data-set  $((x_i, t_i), u_\gamma(x_i, t_i))_{i=1}^m$  with unknown  $\gamma$  by minimizing  $\min_{\gamma \in D} \sum_{i=1}^m (\Phi(\gamma, x_i, t_i) - u_\gamma(x_i, t_i))^2$  via stochastic gradient descent. Moreover, one can obtain estimates for probabilistic quantities related to uncertainty such as

$$\mathbb{V}[u_\Xi(x, t)] \approx \mathbb{V}[\Phi(\Xi, x, t)] \approx \frac{1}{m-1} \sum_{i=1}^m (\Phi(\xi_i, x, t) - \frac{1}{n} \sum_{j=1}^m \Phi(\xi_j, x, t))^2$$

where  $(\xi_i)_{i=1}^m$  are realizations of i.i.d. samples drawn from the distribution of a random variable  $\Xi$  of interest. None of these types of insights were accessible before the presented deep learning method.

We proceed by demonstrating the performance of the proposed algorithm for a general multidimensional affine-linear setting as described in Section 2.1. To this end, let  $d = 3$  and define the initial condition via

$$\varphi_\gamma(x) = \max \left\{ \gamma_\varphi - \frac{1}{3} \sum_{i=1}^3 x_i, 0 \right\}, \quad x \in \mathbb{R}^3,$$

with  $\gamma_\varphi \in D_\varphi \subseteq \mathbb{R}$ . This corresponds to the situation of a Basket put option in a very general multidimensional Black-Scholes model with 3 potentially highly correlated assets. Note that within the above setup, the input dimension of our neural network  $\Phi$  is given by

$$\dim_{\text{in}}(\Phi) = d^3 + 2d^2 + 2d + 1 + 1 = 53.$$

To generate samples of our target variable  $\varphi_\Gamma(S_\Lambda)$ , we simulate solutions of the SDE in (5) using the Euler-Maruyama scheme (8) with  $M = 25$  equidistant steps. Moreover, we use a Monte Carlo

Table 1: Results for the Black-Scholes model

step	avg. time (s)	avg. $\mathcal{L}^1$ -error
0	$0 \pm 0$	$0.6812 \pm 0.0704$
4k	$471 \pm 3$	$0.0088 \pm 0.0056$
8k	$943 \pm 6$	$0.0062 \pm 0.0025$
12k	$1413 \pm 9$	$0.0030 \pm 0.0004$
16k	$1885 \pm 11$	$0.0017 \pm 0.0001$
20k	$2356 \pm 14$	$0.0013 \pm 0.0002$
24k	$2827 \pm 17$	$0.0011 \pm 0.0001$

Table 2: Results for the Basket option

step	avg. time (s)	avg. $\mathcal{L}^1$ -error
0	$0 \pm 0$	$0.7912 \pm 0.0276$
4k	$811 \pm 7$	$0.0131 \pm 0.0019$
8k	$1614 \pm 4$	$0.0087 \pm 0.0013$
12k	$2434 \pm 28$	$0.0062 \pm 0.0009$
16k	$3236 \pm 27$	$0.0058 \pm 0.0011$
20k	$4162 \pm 154$	$0.0046 \pm 0.0007$
24k	$5077 \pm 291$	$0.0042 \pm 0.0002$
28k	$6024 \pm 463$	$0.0039 \pm 0.0001$

approximation with  $m = 2^{20}$  samples to compute the pointwise evaluation of the reference solution  $\bar{u}(\lambda_i)$  according to (10) as needed for the error estimation in (9). Despite the considerably higher dimension of this problem compared with the previous problem from Section 3.1, our deep learning approach shows almost the same approximation behavior, see Table 2. This remarkably weak dependence on the dimension of the input data is further supported by the next examples from physical modelling, where we shall increase the dimensionality of the studied problems even further.

### 3.3 Heat Equation with Varying Diffusion Coefficients

In this Section, we present two examples of high-dimensional heat equations in  $d = 10$  and  $d = 150$  dimensions with paraboloid and Gaussian initial conditions

$$\varphi_\gamma(x) = \|x\|^2 \quad (\text{paraboloid}) \quad \text{and} \quad \varphi_\gamma = e^{-\|x\|^2} \quad (\text{Gaussian}).$$

This formally corresponds to

$$\sigma_\gamma(x) = \gamma_\sigma \quad \text{and} \quad \mu_\gamma(x) = 0$$

where we use a matrix  $\gamma_\sigma \in D_\sigma \subseteq \mathbb{R}^{10 \times 10}$  for the paraboloid case and a scalar  $\gamma_\sigma \in D_\sigma \subseteq \mathbb{R}$  in the Gaussian case, leading to input dimensions of our models  $\Phi$  of

$$\dim_{\text{in}}(\Phi) = d^2 + d + 1 = 111 \quad (\text{paraboloid}) \quad \text{and} \quad \dim_{\text{in}}(\Phi) = d + 1 + 1 = 152 \quad (\text{Gaussian}).$$

Notice that here the solution of the corresponding SDE can be directly sampled via a Brownian motion with random initial position  $X$  and stopping time  $\mathcal{T}$ , i.e.

$$S_\Lambda = X + \sqrt{\mathcal{T}} \Gamma_\sigma N, \quad N \sim \mathcal{N}(0, I_d),$$

see [5, Section 3.2]. For evaluation purposes, these examples were purposefully constructed to have analytic expressions for the parametric solution maps  $\bar{u}$ , which are given by

$$\bar{u}(\gamma_\sigma, x, t) = \|x\|^2 + t \text{Trace}(\gamma_\sigma \gamma_\sigma^*) \quad (\text{paraboloid}), \quad \bar{u}(\gamma_\sigma, x, t) = \frac{e^{-\frac{\|x\|^2}{1+2t\gamma_\sigma^2}}}{(1+2t\gamma_\sigma^2)^{d/2}} \quad (\text{Gaussian}).$$

However, in almost all other practical cases an analytic solution for  $\bar{u}$  is impossible to obtain and numerical methods are the only path forward.

The above dimensionality settings represent regimes which are completely out of scope for all preexisting numerical schemes. Nevertheless, Tables 3 and 4 confirm that our proposed deep learning method once again efficiently converges to the desired parametric solution map  $\bar{u}$ . Our results empirically demonstrate that, contrary to conventional numerical solvers, our deep learning based method does not suffer from the curse of dimensionality, see also Figure 9 in the appendix. We will rigorously prove this fact in the next section.

## 4 Theoretical Guarantees

As a first example, we stick to the heat equation with paraboloid initial condition from above and show that neural networks are capable of simultaneously approximating the parametric solution map  $\bar{u}$  and its gradient with the number of network parameters scaling only polynomially in the dimension  $d$ . Such an approximation guarantee without curse of dimensionality ensures that sensitivity analysis is possible even in very high dimensions. A proof can be found in Section A.1 in the appendix.



Table 3: Results for the heat equation with paraboloid initial condition

step	avg. time (s)	avg. $\mathcal{L}^1$ -error
0	$0 \pm 0$	$0.9609 \pm 0.0052$
4k	$1904 \pm 19$	$0.0150 \pm 0.0008$
8k	$3808 \pm 37$	$0.0120 \pm 0.0007$
12k	$5712 \pm 57$	$0.0093 \pm 0.0006$
16k	$7616 \pm 76$	$0.0068 \pm 0.0001$
20k	$9520 \pm 95$	$0.0062 \pm 0.0003$
24k	$11424 \pm 114$	$0.0057 \pm 0.0001$
28k	$13328 \pm 132$	$0.0056 \pm 0.0000$

Table 4: Results for the heat equation with Gaussian initial condition

step	avg. time (s)	avg. $\mathcal{L}^1$ -error
0	$0 \pm 0$	$0.2035 \pm 0.0714$
4k	$2070 \pm 40$	$0.0123 \pm 0.0047$
8k	$4131 \pm 82$	$0.0050 \pm 0.0018$
12k	$6192 \pm 124$	$0.0051 \pm 0.0022$
16k	$8258 \pm 165$	$0.0033 \pm 0.0015$
20k	$10323 \pm 206$	$0.0025 \pm 0.0011$
24k	$12388 \pm 247$	$0.0024 \pm 0.0008$
28k	$14454 \pm 290$	$0.0019 \pm 0.0002$

**Theorem 3** (Sobolev Approximation). *There exists a neural network  $\Phi$  with ReLU activation function and  $\mathcal{O}(d^4 \log(d/\varepsilon))$  parameters satisfying that*

$$\|\Phi - \bar{u}\|_{\mathcal{L}^\infty(D \times [v,w]^d \times [0,T])} \leq \varepsilon \quad \text{and} \quad \|\nabla\Phi - \nabla\bar{u}\|_{\mathcal{L}^\infty(D \times [v,w]^d \times [0,T])} \leq \varepsilon.$$

Let us now consider the heat equation with varying diffusivity and Gaussian initial condition. In fact, our framework allows us to rigorously prove sample complexity estimates for this problem which represents an almost unique scenario for deep learning based methods. This is rendered possible by the structure of the underlying parametric Kolmogorov PDE and its associated SDE which allows us to describe the distribution of the predictor and target variable, simulate i.i.d. samples, and infer regularity properties on the regression function. We briefly sketch the theorem in the following; the precise formulation and the proof can be found in the appendix.

**Theorem 4** (Generalization). *Using  $s \sim (d/\varepsilon)^2 \text{polylog}(d/\varepsilon)$  many samples, every empirical risk minimizer  $\hat{\Phi}$  of (7) in a suitable hypothesis space  $\mathcal{H}$  of neural networks with ReLU activation function,  $\mathcal{O}(\text{polylog}(d/\varepsilon))$  layers,  $\mathcal{O}(d)$  neurons per layer, and parameters bounded by  $\mathcal{O}(1)$  satisfies with high probability that*

$$\frac{1}{V} \|\hat{\Phi} - \bar{u}\|_{\mathcal{L}^2(D \times [v,w]^d \times [0,T])}^2 \leq \varepsilon$$

where  $V := \text{vol}(D \times [v,w]^d \times [0,T])$ .

Note that it holds that

$$\frac{1}{V} \|\cdot\|_{\mathcal{L}^2(D \times [v,w]^d \times [0,T])}^2 = \|\cdot\|_{\mathcal{L}^2(\mathbb{P}_\Lambda)}^2$$

where  $\mathbb{P}_\Lambda$  is the uniform probability measure on  $D \times [v,w]^d \times [0,T]$ . Thus the estimate in Theorem 4 can be viewed as an estimate in the space  $\mathcal{L}^2(\mathbb{P}_\Lambda)$  and we want to emphasize that our setting easily allows us to choose arbitrary probability measures  $\mathbb{P}$  on  $D \times [v,w]^d \times [0,T]$  and prove analogous results w.r.t. the  $\mathcal{L}^2(\mathbb{P})$  norm.

## 5 Conclusion

The method introduced in this paper is the first deep learning algorithm for the numerical solution of parametric Kolmogorov PDEs and one of few existing algorithms whose use is computationally tractable in high-dimensional settings. The parametric nature of our approach readily allows for sensitivity analysis, model calibration, and uncertainty quantification, all which is of high interest in a variety of applications. Successful numerical experiments in both low- and high-dimensional settings empirically confirm the functionality of the proposed algorithm. In addition, we are able to provide theoretical guarantees for the applicability of our method in high-dimensions.

Besides solving an important problem in scientific computing, our work introduces a class of learning problems that allows for the rigorous investigation of expressivity and sample complexity, along with stable and interpretable algorithms. Such strong results become possible by leveraging the mathematical structure of the learning problem associated with the parametric PDE. We anticipate that the formulation and study of such structured problems will constitute an important future direction of research in the scientific machine learning community as it can enable reliable and interpretable algorithms for the solution of previously intractable problems: in our case parametric families of Kolmogorov PDEs. This contributes substantially to areas like physical modelling of diffusion processes and computational finance, which all rely on the use of such PDEs.

## Broader Impact

The deep-learning technique presented in this work is the first computationally scalable method for the numerical solution of high-dimensional parametric Kolmogorov PDEs. It is also the first method which allows for a straightforward sensitivity analysis of the associated high-dimensional PDE solution manifold with respect to input parameters. In addition, it newly allows for high-dimensional data-driven model calibration and uncertainty quantification. While it is a difficult task to precisely estimate the cascading effects of technological innovations on wider society, it is reasonable to assume that the ubiquity of Kolmogorov equations in science and engineering will lead to a positive impact of our new findings on a multitude of technical areas of social importance.

As an example, Kolmogorov PDEs are heavily used in physics for the modelling of heat flow and diffusion processes [42, 58]. Simultaneously, Fokker-Planck equations, which take the form of Kolmogorov equations in particular special cases, are used in the geophysical and atmospheric sciences as modelling tools for climate change projections [25, 54]. Our described algorithm has clear promise to make previously intractable high-dimensional physical models computationally accessible to scientists. Additionally, our method allows for an easy investigation of changes in complex model forecasts as input parameters are varied during sensitivity analysis. Such advancements have the potential to accelerate scientific research and can directly lead to better predictive models in applied physics and engineering. Reliable and efficient predictive models in turn are essential to rationally inform public policy.

A conceivable risk posed by our work might come in the form of the uncritical use of our algorithm in applications related to financial engineering. The Black-Scholes equation and associated models have been notoriously misused in the last decades by semi-technical users working in financial sectors around the world [31, 59]. The naive usage of technical tools in computational finance has thus likely been a contributing factor to periods of economic instability in recent history. Our technique can now add a powerful solver for high-dimensional parametric PDE problems to the tool kits of individual end-users in finance with various degrees of scientific expertise. Inexperienced users without appropriate quantitative background might be prone to erroneously taking the complexity of a high-dimensional financial model as an indicator for its accuracy. Therefore, one must take great care to systematically inform users without suitable experience in such a scenario that merely increasing the dimension of an inadequate financial model might not necessarily make its results more accurate.

In total, we are confident that the net impact of our work on the scientific community as well as broader society is positive. The probability of uncritical use of our technique and other algorithms in financial engineering can likely be substantially mitigated by targeted educational interventions and we would encourage practical research in this direction. At the same time, we note that our technical contribution is a general-purpose tool which has the potential to stimulate the acceleration of scientific progress in a wide variety of disciplines.

## Acknowledgments and Disclosure of Funding

The research of Julius Berner was supported by the Austrian Science Fund (FWF) under grant I3403-N32. The research of Markus Dablander was supported by the UK EPSRC Centre For Doctoral Training in Industrially Focused Mathematical Modelling (EP/L015803/1).

## References

- [1] C. D. Aliprantis and K. C. Border. Infinite dimensional analysis: a hitchhiker’s guide, 2006.
- [2] W. F. Ames. *Numerical methods for partial differential equations*. Academic press, 2014.
- [3] J. L. Ba, J. R. Kiros, and G. E. Hinton. Layer normalization. *arXiv:1607.06450*, 2016.
- [4] P. Baldi. *Stochastic Calculus: An Introduction Through Theory and Exercises*. Universitext. Springer International Publishing, 2017. ISBN 9783319622262.
- [5] C. Beck, S. Becker, P. Grohs, N. Jaafari, and A. Jentzen. Solving stochastic differential equations and Kolmogorov equations by means of deep learning. *arXiv:1806.00421*, 2018.

- [6] C. Beck, E. Weinan, and A. Jentzen. Machine learning approximation algorithms for high-dimensional fully nonlinear partial differential equations and second-order backward stochastic differential equations. *Journal of Nonlinear Science*, 29(4):1563–1619, 2019.
- [7] J. Berner, D. Elbrächter, P. Grohs, and A. Jentzen. Towards a regularity theory for ReLU networks – chain rule and global error estimates. In *2019 13th International conference on Sampling Theory and Applications (SampTA)*, pages 1–5, 2019.
- [8] 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 Journal on Mathematics of Data Science*, 2(3):631–657, 2020.
- [9] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of political economy*, 81(3):637–654, 1973.
- [10] M. G. Crandall, H. Ishii, and P.-L. Lions. User’s guide to viscosity solutions of second order partial differential equations. *Bulletin of the American mathematical society*, 27(1):1–67, 1992.
- [11] F. Cucker and S. Smale. On the mathematical foundations of learning. *Bulletin of the American mathematical society*, 39(1):1–49, 2002.
- [12] M. Eigel, R. Schneider, P. Trunschke, and S. Wolf. Variational Monte Carlo-bridging concepts of machine learning and high dimensional partial differential equations. *arXiv:1810.01348*, 2018.
- [13] E. Ekström and J. Tysk. The Black–Scholes equation in stochastic volatility models. *Journal of Mathematical Analysis and Applications*, 368(2):498 – 507, 2010.
- [14] D. Elbrächter, P. Grohs, A. Jentzen, and C. Schwab. DNN expression rate analysis of high-dimensional pdes: Application to option pricing. *arXiv:1809.07669*, 2018.
- [15] L. C. Evans and R. F. Gariepy. *Measure Theory and Fine Properties of Functions, Revised Edition*. Textbooks in Mathematics. CRC Press, 2015.
- [16] A. Friedman. *Stochastic Differential Equations and Applications*. Dover Books on Mathematics. Dover Publications, 2012. ISBN 9780486141121.
- [17] J. Gall. *Brownian Motion, Martingales, and Stochastic Calculus*. Graduate Texts in Mathematics. Springer International Publishing, 2016. ISBN 9783319310893.
- [18] M. B. Giles. Multilevel Monte Carlo methods. *Acta Numerica*, 24:259–328, 2015.
- [19] C. Graham and D. Talay. *Stochastic Simulation and Monte Carlo Methods: Mathematical Foundations of Stochastic Simulation*. Stochastic Modelling and Applied Probability. Springer Berlin Heidelberg, 2013. ISBN 9783642393631.
- [20] P. Grohs, F. Hornung, A. Jentzen, and P. Von Wurstemberger. A proof that artificial neural networks overcome the curse of dimensionality in the numerical approximation of Black-Scholes partial differential equations. *arXiv:1809.02362*, 2018.
- [21] P. Grohs, D. Perekrestenko, D. Elbrächter, and H. Bölcskei. Deep Neural Network Approximation Theory. *arxiv:1901.02220*, 2019.
- [22] I. Gühring, G. Kutyniok, and P. Petersen. Error bounds for approximations with deep ReLU neural networks in  $W^{s,p}$  norms. *Analysis and Applications*, pages 1–57, 2019.
- [23] M. Hairer, M. Hutzenthaler, A. Jentzen, et al. Loss of regularity for Kolmogorov equations. *The Annals of Probability*, 43(2):468–527, 2015.
- [24] J. Han, A. Jentzen, and E. Weinan. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- [25] K. Hasselmann. Stochastic climate models part i. theory. *tellus*, 28(6):473–485, 1976.

- [26] 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*, pages 770–778, 2016.
- [27] J. S. Hesthaven and S. Ubbiali. Non-intrusive reduced order modeling of nonlinear problems using neural networks. *Journal of Computational Physics*, 363:55–78, 2018.
- [28] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. ISSN 01621459.
- [29] 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. *arXiv:1901.10854*, 2019.
- [30] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv:1502.03167*, 2015.
- [31] R. A. Jarrow. Risk management models: construction, testing, usage. *The Journal of Derivatives*, 18(4):89–98, 2011.
- [32] A. Jentzen, D. Salimova, and T. Welti. A proof that deep artificial neural networks overcome the curse of dimensionality in the numerical approximation of Kolmogorov partial differential equations with constant diffusion and nonlinear drift coefficients. *arXiv:1809.07321*, 2018.
- [33] Y. Khoo, J. Lu, and L. Ying. Solving parametric PDE problems with artificial neural networks. *arXiv:1707.03351*, 2017.
- [34] P. E. Kloeden and E. Platen. *Numerical solution of stochastic differential equations*, volume 23 of *Applications of Mathematics (New York)*. Springer-Verlag, Berlin, 1992. ISBN 3-540-54062-8.
- [35] P. E. Kloeden and E. Platen. *Numerical solution of stochastic differential equations*, volume 23. Springer Science & Business Media, 2013.
- [36] G. Kutyniok, P. Petersen, M. Raslan, and R. Schneider. A theoretical analysis of deep neural networks and parametric pdes. *arXiv:1904.00377*, 2019.
- [37] F. Laakmann and P. Petersen. Efficient approximation of solutions of parametric linear transport equations by ReLU DNNs. *arXiv:2001.11441*, 2020.
- [38] L. Li, K. Jamieson, A. Rostamizadeh, E. Gonina, M. Hardt, B. Recht, and A. Talwalkar. Massively parallel hyperparameter tuning. *arXiv:1810.05934*, 2018.
- [39] R. Liaw, E. Liang, R. Nishihara, P. Moritz, J. E. Gonzalez, and I. Stoica. Tune: A research platform for distributed model selection and training. *arXiv:1807.05118*, 2018.
- [40] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. *arXiv:1711.05101*, 2017.
- [41] G. F. Montufar, R. Pascanu, K. Cho, and Y. Bengio. On the number of linear regions of deep neural networks. In *Advances in neural information processing systems*, pages 2924–2932, 2014.
- [42] A. Pascucci. Kolmogorov equations in physics and in finance. In *Elliptic and parabolic problems*, pages 353–364. Springer, 2005.
- [43] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems* 32, pages 8024–8035. Curran Associates, Inc., 2019.
- [44] P. Petersen and F. Voigtlaender. Optimal approximation of piecewise smooth functions using deep ReLU neural networks. *Neural Networks*, 108:296–330, 2018.

- [45] O. Pironneau and Y. Achdou. Partial differential equations for option pricing. *Handbook of Numerical Analysis*, 15:369–495, 2009.
- [46] D. Pollard. *A User’s Guide to Measure Theoretic Probability*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2002. ISBN 9780521002899.
- [47] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [48] C. Reisinger and Y. Zhang. Rectified deep neural networks overcome the curse of dimensionality for nonsmooth value functions in zero-sum games of nonlinear stiff systems. *arXiv:1903.06652*, 2019.
- [49] S. Ruder. An overview of gradient descent optimization algorithms. *arXiv:1609.04747*, 2016.
- [50] C. Schwab and J. Zech. Deep learning in high dimension: Neural network expression rates for generalized polynomial chaos expansions in UQ. *Analysis and Applications*, 17(01):19–55, 2019.
- [51] R. Seydel. *Tools for computational finance*, volume 3. Springer, 2006.
- [52] R. Seydel. *Tools for Computational Finance*. Universitext. Springer Berlin Heidelberg, 2013. ISBN 9783662047118.
- [53] J. Sirignano and K. Spiliopoulos. DGM: A deep learning algorithm for solving partial differential equations. *Journal of Computational Physics*, 375:1339–1364, 2018.
- [54] J. Thuburn. Climate sensitivities via a fokker–planck adjoint approach. *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography*, 131(605):73–92, 2005.
- [55] V. Vapnik. *Statistical learning theory. 1998*, volume 3. Wiley, New York, 1998.
- [56] E. Weinan and B. Yu. The deep Ritz method: a deep learning-based numerical algorithm for solving variational problems. *Communications in Mathematics and Statistics*, 6(1):1–12, 2018.
- [57] E. Weinan, J. Han, and A. Jentzen. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Communications in Mathematics and Statistics*, 5(4):349–380, 2017.
- [58] D. V. Widder. *The heat equation*, volume 67. Academic Press, 1976.
- [59] P. Wilmott. The use, misuse and abuse of mathematics in finance. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 358(1765):63–73, 2000.
- [60] D. Yarotsky. Error bounds for approximations with deep ReLU networks. *Neural Networks*, 94: 103–114, 2017.
- [61] F. Yu, D. Wang, E. Shelhamer, and T. Darrell. Deep layer aggregation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2403–2412, 2018.