

Solving parametric PDE problems with artificial neural networks

Yuehaw Khoo* Jianfeng Lu † Lexing Ying‡

July 14, 2017

Abstract

The curse of dimensionality is commonly encountered in numerical partial differential equations (PDE), especially when uncertainties have to be modeled into the equations as random coefficients. However, very often the variability of physical quantities derived from PDE can be captured by a few features on the space of random coefficients. Based on such observation, we propose using neural-network, a technique gaining prominence in machine learning tasks, to parameterize the physical quantity of interest as a function of random input coefficients. The simplicity and accuracy of the approach are demonstrated through notable examples of PDEs in engineering and physics.

1 Introduction

Uncertainty quantifications in physical and engineering applications often involve the study of partial differential equations (PDE) with random coefficients. To understand the behavior of a system in the presence of uncertainties, one can extract PDE derived physical quantities as functionals of the random fields. This can potentially require solving the PDE exponential number of times numerically even with suitable discretization of the PDE domain, and of the range of random variables. Fortunately in PDE applications, often these functionals depend only on a few characteristic “features” of the random fields, allowing them to be determined from solving the PDE limited number of times.

A commonly used approach for uncertainty quantifications is Monte-Carlo sampling. An ensemble of solutions is built by repeatedly solving the PDE with different realizations. Then physical quantities of interest, for example, the mean of the solution at a given location, can be computed from the ensemble of

*Department of Mathematics, Stanford University, Stanford, CA 94305, USA (ykhoo@stanford.edu).

†Department of Mathematics, Department of Chemistry and Department of Physics, Duke University, Durham, NC 27708, USA (jianfeng@math.duke.edu).

‡Department of Mathematics and ICME, Stanford University, Stanford, CA 94305, USA (lexing@stanford.edu).

solutions. Although being applicable in many situations, the computed quantity is inherently noisy. Moreover, it lacks the ability to obtain new solutions if they are not sampled previously. Other approaches exploit the low underlying dimensionality assumption in a more direct manner. For example the stochastic Galerkin method [12, 15] expands the random solution using certain prefixed basis functions (i.e. polynomial chaos [18, 19]) on the space of random variables, thereby reducing the high dimensional problem to a few deterministic PDEs. Such type of methods requires careful treatment of the uncertainty distributions, and since the basis used is problem independent, when the dimensionality of the random variables is high the method could be expensive. There are data driven approaches for basis learning such as applying Karhunen-Loève expansion to PDE solutions from different realizations of the PDE [3]. Similar to the related principal component analysis, such linear dimension-reduction technique may not fully exploit the nonlinear interplay between the random variables. At the end of day, the problem of uncertainty quantification is one of characterizing the low-dimensional structure of the random field that gives the observed quantities.

On the other hand, the problem of dimensionality reduction has been central to the fields of statistics and machine learning. The fundamental task of regression seeks to find a function h_θ parameterized by parameter vector $\theta \in \mathbb{R}^p$ such that

$$f(a) \approx h_\theta(a), \quad a \in \mathbb{R}^n. \quad (1)$$

However, choosing a sufficiently large class of approximation functions without the issue of over-fitting remains a delicate business. As an example, in linear regression, the standard procedure is to fix a set of basis (or feature maps) $\{\phi_k(a)\}$ such that

$$f(a) = \sum_k \beta_k \phi_k(a) \quad (2)$$

and determine the parameter β_k 's from sampled data. The choice of basis is important to the quality of regression, just as in the case of studying PDE with random coefficients. Recently, the deep neural-network shows unprecedented success in solving a variety of difficult regression problems related to pattern recognitions [8, 11, 14]. A key advantage of neural-network is that it bypasses the traditional need to handcraft basis for spanning $f(a)$ but instead, learns the optimal basis that satisfies (1) directly from data. The performance of neural-network in machine learning applications, and more recently in physical applications such as representing quantum many-body states (e.g. [16, 2]), prompts us to study its use in the context of solving PDE with random coefficients. More precisely, we want to learn $f(a)$ that maps the random coefficient vectors a in a PDE to some physical quantity described by the PDE.

The idea of representing a function by neural-network to perform model reduction may seem straight-forward to machine learning practitioners. However, to the best of our knowledge, such technique has not been adapted to solving PDE with uncertainties. Through the demonstrations on two simple examples, we hope this short paper could open a path to uncertainty quantification with less pain. We note that our work is substantially different from [10, 13] which

solve deterministic PDE numerically using a neural-network. The goal of these works is to parameterize the solution of a deterministic PDE using neural-network and replace Galerkin-type methods when performing model reduction. It is also different from [7] where a deterministic PDE is solved as a stochastic control problem using neural-network. In this paper, the function we want to parameterize is over the coefficient field of the PDE.

The advantages of having an explicitly parameterized approximation to $f(a)$ are numerous, which we will only list a couple here. First, the neural-network parameterized function can serve as a surrogate forward model for generating samples cheaply for statistical analysis. Second, the task of optimizing some function of the physical quantity with respect to the PDE coefficients can be done with the help of a gradient calculated from the neural-network. To summarize, obtaining a neural-network parametrization could limit the use of expensive PDE solvers in applications.

We demonstrate the success of neural-network in two PDE applications. In Section 2, we describe the overall approach for handling these PDE problems. In Section 3, we solve the effective conductance in non-homogeneous media where the local conductivities are random and the ground state energy of a nonlinear Schrödinger equation (NLSE) having random potential. In Section 4, we provide a few notes on the implementation details and we finally conclude in Section 5.

2 Overall framework

Our approach to solving quantities arise from PDE with randomness is embarrassingly simple, consisting of the following steps:

- (1) Sample the random coefficients (a in (1)) of the PDE from a user-specified distribution. For each set of coefficients, solve the deterministic PDE to obtain the physical quantity of interest ($f(a)$ in (1)).
- (2) Use a neural-network as the surrogate model $h_\theta(a)$ in (1) and train it using the previously obtained samples.
- (3) Validate the surrogate forward model with more samples. The neural network is now ready for applications.

In Section 2.1, we introduce some standard neural-network terminologies. For further details, we refer readers to [17, 6]. Readers who are already familiar with neural-network may skip this section. In Section 2.2 we describe the neural-network architecture we use in our numerical experiments. As we shall see, the neural-network we use is rather simple compared to the typical architectures used in machine learning. Our numerical examples will show that even with such simple architecture, the neural-network gives good performance for the particular problems we have. PDE domain knowledge would certainly help proposing better neural-network architecture to improve the performance, which can be explored for specific parametric PDE problems.

2.1 Neural-network operations

Neural-network can be pictured as a sequence of layers with non-linear mapping between two adjacent layers. The purpose of training is to determine the parameters of these nonlinear maps. The element of the final layer is used as $f(a)$. We mainly use three types of layers to build an artificial neural-network: convolutional layer, pooling layer and fully/densely-connected layer. In this paper, we represent a layer as a 3-tensor with size $\alpha \times L_x \times L_y$ (see Fig. 1 for example) where the first dimension is called the *channel dimension*, and the second and third dimensions are usually called the *spatial dimensions*. In all figures, the dimensions with size $L_x \times L_y$ lie within the plane of paper and the dimension of size α extends out of paper.

Fig. 1 illustrates the operation that maps layer B to a convolutional layer C . Denote the elements of B within the tube with dashed edges (size $\alpha \times K_x \times K_y$) as $v^{i,j} \in \mathbb{R}^{\alpha K_x K_y}$ where (i, j) indicates the center of the tube along the spatial-dimensions. Here $K_x \times K_y$ is called the *kernel window size*. An element of the convolution layer C is obtained via the map

$$C_{kij} = s((W^k)^\top v^{i,j} + b^k), \quad W^k \in \mathbb{R}^{\alpha K_x K_y}, \quad b^k \in \mathbb{R} \quad (3)$$

where s is a pre-specified *nonlinearity* (nonlinear function), usually chosen from the rectified-linear unit (ReLU) or sigmoid function. It should be clear that the *filter* W^k , and also b^k remain unchanged when being applied to different spatial locations (i, j) in B . Therefore the map between two layers are parameterized by $\alpha_2 \alpha_1 K_x K_y + \alpha_2$ parameters where α_1, α_2 are channel numbers of B, C respectively. The convolutional map can be understood as applying a nonlinear filter to a spatial neighborhood to extract its important features. It is clear from Fig. 1 that only the elements towards the spatial center of B have their spatial neighbors properly covered by the filters. As a consequence, the convolutional map reduces the sizes of spatial dimensions. Therefore often certain *padding* operation is required to enlarge the input-layer B in order for the filters to properly cover the locations near the edges of the original input.

Another important operation is pooling. A pooling operation decimates a layer B along the spatial dimensions to obtain a pooling layer C . B, C have the same channel number, while the sizes of spatial dimensions of C are reduced compare to B . In this paper, we only use *sum-pooling* which sums up B across the spatial dimensions:

$$C_k = \sum_{i,j} B_{kij}. \quad (4)$$

Therefore in this case C is a vector with size being the channel number of B . This operation does not introduce any extra parameter.

Lastly, we introduce the densely-connected layer. A densely-connected layer C is obtained from input B via the map

$$C(\cdot) = s(WB(\cdot) + b) \quad (5)$$

parameterized by matrix W and vector b with appropriate sizes, where $B(\cdot), C(\cdot)$ denote the vectorization of B, C . Here s is an element-wise nonlinearity.




Figure 1 Example of a convolutional map between layers B and C with sizes $5 \times 11 \times 11$ and $3 \times 8 \times 8$ respectively. These sizes are indicated by the triplet of numbers “ $\alpha @ L_x \times L_y$ ” in the diagram in order to make clear the unequal footing between the channel and spatial dimensions. In the figure, $L_x \times L_y$ is the size of the 2D squares lying within the plane of paper and α is the total number of squares.

The parameters of the inter-layer mappings are optimized to satisfy (1). In this paper, we choose a least-squares penalty to enforce (1).

2.2 Proposed network architecture

We introduce the main workhorse in this paper - artificial neural-network with a single convolutional layer in Fig. 2, for solving PDE with periodic boundary conditions. We present the architecture for the 2D case with domain being a unit square, though it can be generalized to solving PDEs in any dimensions. The input is a random matrix $a \in \mathbb{R}^{n \times n}$ resembling the random field on grid points, and the output of the network gives physical quantity of interest from the PDE. The main part of the network is a convolutional layer with ReLU being the nonlinearity. This extracts the relevant features of the random field around each grid-point that contribute to the final output. Finally, a densely-connected layer without nonlinearity is used to create the output. All operations in Fig. 2 are standard except the padding operation. Typically, zero-padding is used to enlarge the size of the input in image classification task, whereas we extend the input periodically due to the assumed boundary condition.

The motivation of the architecture is straightforward. As in (2), we look for an approximation to function $f(a)$ via the choice of basis $\{\tilde{\phi}_{ijk}(a)\}$ and weights $\{\tilde{\beta}_{ijk}\}$ such that

$$f(a) \approx \sum_{k=1}^{\alpha} \sum_{i=1}^n \sum_{j=1}^n \tilde{\beta}_{kij} \tilde{\phi}_{kij}(a) \quad (6)$$

for a drawn from a certain distribution. This exactly corresponds to the operations in Fig. 2. First $\tilde{\phi}_{kij}(a)$'s are obtained via a convolutional layer (each $\tilde{\phi}_{kij} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$ is the composition of a nonlinearity and convolution), then




Figure 2 Single convolutional layer neural network. The title between two adjacent layers indicates the neural-network operation between them, which also contains the type of nonlinearity used. The triplet of numbers “ $\alpha @ L_x \times L_y$ ” indicates the channel number and the sizes of spatial dimensions.

$\tilde{\phi}_{kij}(a)$ ’s are linearly-combined via densely-connected layer.

2.2.1 Incorporating symmetry

The reason for using a convolutional neural-network with same set of filters covering each input’s location is that in this paper, the function $f(a)$ to be approximated is shift invariant. Therefore one naturally expects the same set of features being extracted when the random field a is translated spatially. More precisely, let $a_{ij}^{\tau_1\tau_2} := a_{(i+\tau_1)(j+\tau_2)}$ where the additions are done on \mathbb{Z}_n . The use of convolutional layer implies

$$\tilde{\phi}_{kij}(a^{\tau_1\tau_2}) = \tilde{\phi}_{k(i-\tau_1)(j-\tau_2)}(a), \quad \forall(\tau_1, \tau_2) \in \{1, \dots, n\}^2. \quad (7)$$

Another point is that when the function considered satisfies

$$f(a) = f(a^{\tau_1\tau_2}), \quad \forall(\tau_1, \tau_2) \in \{1, \dots, n\}^2, \quad (8)$$

the number of parameters can be further reduced. In the ideal case where

$$f(a) = \sum_{k=1}^{\alpha} \sum_{i=1}^n \sum_{j=1}^n \tilde{\beta}_{kij} \tilde{\phi}_{kij}(a) \quad (9)$$

for some 3-tensor $\tilde{\beta}$, translational invariance (8) implies

$$\begin{aligned}
 f(a) &= \sum_{k=1}^{\alpha} \sum_{i=1}^n \sum_{j=1}^n \tilde{\beta}_{kij} \tilde{\phi}_{kij}(a^{\tau_1\tau_2}) \\
 &= \sum_{k=1}^{\alpha} \sum_{i=1}^n \sum_{j=1}^n \tilde{\beta}_{kij} \tilde{\phi}_{k(i-\tau_1)(j-\tau_2)}(a) \quad \forall(\tau_1, \tau_2) \\
 &= \frac{1}{n^2} \sum_{(\tau_1, \tau_2)} \sum_{k=1}^{\alpha} \sum_{i=1}^n \sum_{j=1}^n \tilde{\beta}_{kij} \tilde{\phi}_{k(i-\tau_1)(j-\tau_2)}(a)
 \end{aligned}$$




Figure 3 Single convolutional layer neural network for representing translational invariant function.

$$\begin{aligned}
&= \sum_{k=1}^{\alpha} \left(\frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \tilde{\beta}_{kij} \right) \left(\sum_{(\tau_1, \tau_2)} \tilde{\phi}_{k(i-\tau_1)(j-\tau_2)}(a) \right) \\
&= \sum_{k=1}^{\alpha} \beta_k \phi_k(a)
\end{aligned} \tag{10}$$

where $\beta_k := \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \tilde{\beta}_{kij}$ and $\phi_k(a) := \sum_{i=1}^n \sum_{j=1}^n \tilde{\phi}_{kij}(a)$. The second equality follows from (7). Eq. (10) implies instead of mapping the convolution layer with entries being $\tilde{\phi}_{kij}(a)$'s to a final output using the 3-tensor $\tilde{\beta} \in \mathbb{R}^{\alpha \times n \times n}$, without loss of generality, one can simply use a vector $\beta \in \mathbb{R}^{\alpha}$ after a sum-pooling of the convolutional layer. This leads to the architecture in Fig. 3 for representing a translational invariant function with less parameters.

3 Numerical examples

In this section, using the neural-network architecture presented in Section 2.2, we approximate physical quantities in two PDE applications as a function of random coefficients.

3.1 Effective coefficients

Take the simplest example of finding the effective conductance in a non-homogeneous media with random conductivities. For this, we consider the elliptic equation

$$\nabla \cdot a(x)(\nabla u(x) + \xi) = 0, \quad x \in [0, 1]^d \tag{11}$$

with periodic boundary condition where $\xi \in \mathbb{R}^d, |\xi|^2 = 1$ ($|\cdot|$ is the Euclidean norm). To ensure ellipticity, we consider the class of coefficient functions

$$\mathcal{A} = \{a \in L^\infty([0, 1]^d) \mid 1 \geq a \geq \lambda_0 > 0\}$$

For a given ξ , we want to obtain the effective conductance functional $A_{\text{eff}} : \mathcal{A} \rightarrow \mathbb{R}$ defined by

$$A_{\text{eff}}(a(\cdot)) = \int_{[0, 1]^d} a(x) |\nabla u(x) + \xi|^2 dx. \tag{12}$$

To solve the PDE and obtain training samples, we discretize the domain into $x_i = i/n$ where the multi-index $i \in \{(i_1, \dots, i_d)\}$, $i_1, \dots, i_d = 1, \dots, n$. Then the action of Laplace operator on u is discretized as

$$\sum_{k=1}^d a_{i+\frac{1}{2}e_k} (u_{i+e_k} - u_i) + a_{i-\frac{1}{2}e_k} (u_{i-e_k} - u_i) \quad (13)$$

for each i , where $\{e_k\}_{k=1}^d$ denotes the canonical basis in \mathbb{R}^d . Here $a_i := a(i/n)$ and $u_i := u(i/n)$. $\nabla \cdot a(x)\xi$ is discretized using central difference. For simplicity, we make the further approximation $a_{i+e_k/2} = (a_{i+1} + a_i)/2$. After such approximations, a becomes a vector in \mathbb{R}^{n^d} . We assume $\{a_i\}_{i=1}^{n^d}$ are independently and identically distributed according to $\mathcal{U}[0.3, 3]$ where $\mathcal{U}[l_1, l_2]$ denotes the uniform distribution on the interval $[l_1, l_2]$.

The results of learning the effective conductance function are presented in Table 1. We use the same number of samples for training and validation. Both the training and validation error are measured by

$$\sqrt{\frac{\sum_k (h_\theta(a^k) - A_{\text{eff}}(a^k))^2}{\sum_k A_{\text{eff}}(a^k)^2}}, \quad i = 1, \dots, n^d \quad (14)$$

where a^k 's can either be the training or validation samples sampled from the same distribution and h_θ is the neural-network-parameterized approximation function. We note that the choice of small grid size is not due to the computational cost of network training as these days, deep neural-networks for image classification typically have at least 256×256 input size and much more complicated architecture. For the diffusion problem with i.i.d. random coefficients, the effective conductance $A_{\text{eff}}(a)$ concentrates with a standard deviation that scales like $1/\sqrt{n^d}$ [5]. Therefore, the problem of learning $A_{\text{eff}}(a)$ is asymptotically easier when the grid size gets larger.

n	α	Training error	Validation error	Average A_{eff}	No. of samples	No. of parameters
8	16	2.4×10^{-3}	3.0×10^{-3}	1.86 ± 0.10	1.2×10^4	1057
16	16	2.1×10^{-3}	2.2×10^{-3}	1.87 ± 0.052	2.4×10^4	4129

Table 1 Error in approximating the effective conductance function $A_{\text{eff}}(a)$ in 2D. The mean and standard deviation of the effective conductance are computed from the samples in order to show the variability. The sample sizes for training and validation are the same.

Before concluding this subsection, we use the exercise of determining the effective conductance in 1D to further motivate the usage of a neural-network. In 1D the effective conductance can be expressed analytically as the harmonic mean of a_i 's:

$$A_{\text{eff}}(a) = \left(\frac{1}{n} \sum_{i=1}^n \frac{1}{a_i} \right)^{-1}. \quad (15)$$




Figure 4 Neural-network architecture for approximating $A_{\text{eff}}(a)$ in the 1D case. Although the layers in third stage are essentially densely-connected layers, we still identify them as convolution layers to reflect the symmetry between the first and third stages.




Figure 5 The first stage's output of the neural-network in Fig. 4 fitted by $\beta_1/x + \beta_2$.

This function indeed approximately corresponds to the deep neural-network shown in Fig. 4. The neural-network is separated into three stages. In the first stage, the approximation to function $1/a_i$ is constructed for each a_i by applying a few convolution layers with size 1 kernel window. In this stage, the channel size for these convolution layers is chosen to be 16 except the last layer since the output of the first stage should be a vector of size n . In the second stage, a layer of sum-pooling with size n window is used to perform the summation in (15), giving a scalar output. The third and first stages have the exact same architecture except the input to the third stage is a scalar. 2560 samples are used for training and another 2560 samples are used for validation. We let $a_i \sim \mathcal{U}[0.3, 1.5]$, giving an effective conductance of 0.77 ± 0.13 for $n = 8$. We obtain 4.9×10^{-4} validation error with the neural-network in Fig. 4 while with the network in Fig. 3, we get 5.5×10^{-3} accuracy with $\alpha = 16$. As a check, in Fig. 5 we show that the output from the first stage is well-fitted by the reciprocal function.

We remark that although incorporating domain knowledge to build a sophisticated neural-network architecture would likely boost the approximation quality, even a simple network as in Fig. 3 can already give decent results.

3.2 NLSE with random potential

In the second example, we want to find the ground state energy E_0 of a nonlinear Schrödinger equation (NLSE) with random potential

$$-\nabla^2 u(x) + V(x)u(x) + \sigma u(x)^3 = E_0 u(x), \quad x \in [0, 1]^2 \quad (16)$$

where $\sigma = 2$ and $V(x)$ a random field, subject to the normalization constraint

$$\|u\|_{L^2([0,1]^2)} = 1. \quad (17)$$

This is a defocusing cubic Schrödinger equation, which can be understood as a model for soliton in nonlinear photonics or Bose-Einstein condensate with random media. It is more difficult to solve for the NLSE numerically compare to (11). Therefore in this case, the value of having a surrogate model of E_0 as a functional of $V(\cdot)$ is more significant.

To solve the problem numerically, the domain is discretized similarly as in Section 3.1 and the Laplace operator is discretized using standard 5-point stencil. Now the goal is to obtain a neural-network parametrization for $E_0(\cdot)$, with input now being $V \in \mathbb{R}^{n^2}$, the discretization of the random field $V(x)$, with i.i.d. entries distributed according to $\mathcal{U}[1, 16]$. In order to generate training samples, for each realization of V , the nonlinear eigenvalue problem (16) subject to the normalization constraint (17) is solved by homotopy method. First, the case $\sigma = 0$ is solved as a standard eigenvalue problem. Then σ is changed from 0 to 2 with 0.4 step size. For each σ , Newton’s method is used to solved the NLSE for $u(x, V)$ and $E_0(V)$, using $u(x, V)$ and $E_0(V)$ corresponding to the previous σ as initialization. The results are presented in Table 2.

n	α	Training error	Validation error	Average E_0	No. of samples	No. of parameters
8	5	4.9×10^{-4}	5.0×10^{-4}	10.48 ± 0.51	4800	331
16	5	1.5×10^{-4}	1.5×10^{-4}	10.46 ± 0.27	1.05×10^4	1291

Table 2 Error in approximating the lowest energy level $E_0(V)$ for $n = 8, 16$ discretization.

4 Notes on implementation

The neural-network is implemented using Keras [4], an application programming interface running on top of TensorFlow [1] (a library of toolboxes for training neural-network). We use a mean-squared-error loss function. The optimization

is done using a variant of stochastic gradient descent, the Adam optimizer [9]. The hyper-parameter we tune is the learning rate, which we lower if the training error fluctuates too much. The weights are initialized randomly from the normal distribution. The input to the neural-network is whitened to have unit variance and zero-mean on each dimension. The mini-batch size is always set to between 50 and 200.

5 Conclusion

In this note, we present method based on deep neural-network to solve PDE with random coefficients. Physical quantities of interest are learned as a function of the random coefficients. The numerical experiments on diffusion equation and NLSE show the effectiveness of simple feed-forward neural network in parametrizing such function to 10^{-3} accuracy. We remark that while many questions should be asked, such as what is the best network architecture and what situations can this approach handle, the goal of this short note is simply to suggest neural-network as a promising tool for model reduction when solving PDE with uncertainties.

References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Josh Kudlur, Manjunath Levenberg, Dan Mane, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viegas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng, *Tensorflow: Large-scale machine learning on heterogeneous distributed systems*, arXiv preprint arXiv:1603.04467 (2016).
- [2] Giuseppe Carleo and Matthias Troyer, *Solving the quantum many-body problem with artificial neural networks*, *Science* **355** (2017), no. 6325, 602–606.
- [3] Mulin Cheng, Thomas Y Hou, Mike Yan, and Zhiwen Zhang, *A data-driven stochastic method for elliptic PDEs with random coefficients*, *SIAM/ASA Journal on Uncertainty Quantification* **1** (2013), no. 1, 452–493.
- [4] François Chollet, *Keras (2015)*, URL <http://keras.io> (2017).
- [5] Antoine Gloria and Felix Otto, *An optimal variance estimate in stochastic homogenization of discrete elliptic equations*, *The annals of probability* **39** (2011), no. 3, 779–856.

- [6] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep learning*, MIT press, 2016.
- [7] Jiequn Han, Arnulf Jentzen, and Weinan E, *Overcoming the curse of dimensionality: Solving high-dimensional partial differential equations using deep learning*, arXiv preprint arXiv:1707.02568 (2017).
- [8] Geoffrey E Hinton and Ruslan R Salakhutdinov, *Reducing the dimensionality of data with neural networks*, Science **313** (2006), no. 5786, 504–507.
- [9] Diederik Kingma and Jimmy Ba, *Adam: A method for stochastic optimization*, arXiv preprint arXiv:1412.6980 (2014).
- [10] Isaac E Lagaris, Aristidis Likas, and Dimitrios I Fotiadis, *Artificial neural networks for solving ordinary and partial differential equations*, IEEE Transactions on Neural Networks **9** (1998), no. 5, 987–1000.
- [11] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, *Deep learning*, Nature **521** (2015), no. 7553, 436–444.
- [12] Hermann G Matthies and Andreas Keese, *Galerkin methods for linear and nonlinear elliptic stochastic partial differential equations*, Computer methods in applied mechanics and engineering **194** (2005), no. 12, 1295–1331.
- [13] Keith Rudd and Silvia Ferrari, *A constrained integration (CINT) approach to solving partial differential equations using artificial neural networks*, Neurocomputing **155** (2015), 277–285.
- [14] Jürgen Schmidhuber, *Deep learning in neural networks: An overview*, Neural networks **61** (2015), 85–117.
- [15] George Stefanou, *The stochastic finite element method: past, present and future*, Computer Methods in Applied Mechanics and Engineering **198** (2009), no. 9, 1031–1051.
- [16] Giacomo Torlai and Roger G Melko, *Learning thermodynamics with Boltzmann machines*, Physical Review B **94** (2016), no. 16, 165134.
- [17] Stanford University, *CS231n convolutional neural networks for visual recognition (2015)*, URL <http://cs231n.github.io/> (2017).
- [18] Norbert Wiener, *The homogeneous chaos*, American Journal of Mathematics **60** (1938), no. 4, 897–936.
- [19] Dongbin Xiu and George Em Karniadakis, *The Wiener–Askey polynomial chaos for stochastic differential equations*, SIAM journal on scientific computing **24** (2002), no. 2, 619–644.