Spectral PINNs: Fast uncertainty propagation with physics-informed neural networks

Anonymous Author(s)
Affiliation
Address
email

Abstract

Physics-informed neural networks (PINNs) promise to significantly speed up partial differential equation (PDE) solvers. However, most PINNs can only solve deterministic PDEs. Here, we consider *stochastic* PDEs that contain partially unknown parameters. Specifically, we aim to perform fast uncertainty propagation, i.e., quantifying uncertainties in the PDE-solving function if parameters are unknown. Classical uncertainty quantification methods such as Monte Carlo sampling, stochastic Galerkin, collocation, or discrete projection methods become computationally too expensive with an increasing number of stochastic parameters. Here, we combine neural networks with a well-known technique from computational fluid dynamics called spectral expansion or polynomial chaos expansion. The success of spectral expansions can be explained by the achieved separation of spatiotemporal and spectral domain, theoretical guarantees, and fast computation of stochastic summaries (e.g., mean), but is reduced by the computational cost. Our *Spectral PINNs* approximate the underlying spectral coefficients with a neural network and reduce the computational cost of the spectral expansion while maintaining guarantees. We derive the method for partial differential equations, discuss runtime, demonstrate initial results on the convection-diffusion equation, and provide steps towards convergence guarantees.

1 Introduction

Solving partial differential equations (PDEs) in fluid dynamics, astrophysics, or climate science is computationally expensive [1]. For example, solving a climate model at 1km resolution can take 10 days on a 5000-GPU supercomputer [2]. Decades of research in numerical methods [3] and physical approximations [4] has achieved drastic reductions in the complexity of solving PDEs. However, physical approximations require extensive domain knowledge and often sacrifice accuracy [4]. Numerical methods remain computationally expensive for high-dimensional problems or uncertainty quantification [5].

Recently, researchers in the field of physics-informed neural networks (PINNs) have proposed to combine numerical methods with deep neural networks (NNs) [6, 5]. The combination builds a perfect team: NNs are capable of high-dimensional function approximation and fast inference times, but lack guarantees [7]. Numerical methods are slow for high-dimensional problems, but have rigorous convergence guarantees and physical constraints [8]. Each balancing the other's weakness, physical knowledge has been integrated into NNs as inputs [9, 10], loss function [6], latent state [11, 12, 13, 14], analytic constraints [15, 16], or evaluation function [17, 18]. From the other side, NNs have been embedded in physical PDEs as parameters [19, 6], dynamics [20], residual [21, 22], differential operator [23, 24], or solution [6, 25]. Our *Spectral PINNs* leverage NNs

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

to approximate the solution of a PDE, creating a learning-based surrogate model [11] that can be used for fast propagation of uncertainties or sensitivity analyses.

Despite the significant success of PINNs to speed-up PDE solvers [25], most methods are limited to deterministic PDEs [5]. In practice, however, many parameters of PDEs are partially unknown and uncertainty quantification plays a vital role in the development of real-life safety-critical systems [26], climate risk estimation [27], or parameter inference [28]. A small set of PINNs has targeted stochastic differential equations [29, 30, 31], but has been developed rather independently from the existing tools in uncertainty quantification (UQ) [28]. Our work builds a bridge between UQ and deep learning by combining well-known spectral expansions with neural networks.

Uncertainty propagation, the UQ subfield of our interest, aims to calculate statistical summaries of the solution given knowledge about parameter distributions [28]. We can roughly divide methods for uncertainty propagation in intrusive and non-intrusive techniques. Intrusive techniques, such as Stochastic Galerkin [32], require domain knowledge of the PDEs and modifiability of the PDE solver which often does not exist [8]. Non-intrusive techniques, such as Monte Carlo sampling [33], collocation methods [34], or discrete projection [34], are purely data-driven methods that only assume solver-generated datasets. The runtime complexity of many non-intrusive techniques, however, can scale cubically with the number of stochastic parameters [32]. Deep learning-based generative models can be fast, but often come without theoretical guarantees [35, 36, 37, 38]. Our method is non-intrusive, faster than conventional UQ methods, and leverages a widely adopted UQ technique called spectral expansion or polynomial chaos expansion (PCE) [39, 40, 28].

In particular, we propose to approximate the spectral coefficients with NNs. The integration of NNs into spectral expansions enables fast calculation of statistical summaries, sensitivity analyses [41], and the potential for theoretical guarantees of optimality or convergence [28]. The closest to our work proposes NNs to estimate the coefficients of an arbitrary polynomial chaos [42], but has only been demonstrated on a 1D stochastic elliptic diffusion equation, not on PDEs, and is intrusive [43].

We contribute *Spectral PINNs*, specifically:

- The extension of spectral PINNs to PDEs, spatiotemporal, parabolic, hyperpolic, and evolutionary equations.
- A scalable, mesh-free, and data-driven method for non-intrusive uncertainty propagation in high-dimensional parameter and spatiotemporal spaces.
- A potentially faster method for uncertainty propagation than traditional UQ methods, such as stochastic Galerkin or collocation.
- The first demonstration of *Spectral PINNs* on the convection-diffusion equation.

2 Approach

Quantifying uncertainties in PDE solutions via computing large ensembles is computationally expensive. We approximate the coefficients of a spectral expansion with NNs to build a fast and accurate surrogate for uncertainty propagation.

2.1 Problem definition

First, we define a stochastic PDE with the solution, u, as:

$$\mathcal{N}_x[u(t,x;\omega);\boldsymbol{Q}(x;\omega)] = 0,\ \forall\ t{\in}D_t,\ x{\in}D_x,\ \omega{\in}\Omega,$$
$$\mathcal{B}_{t,x}[u(t,x;\omega)] = 0,\ \forall\ t{\in}\partial D_t,\ x{\in}\partial D_x,$$

with spatial domain, D_x, temporal domain, D_t, sample space, Ω, domain boundary, ∂, potentially nonlinear operator, $\mathcal{N}$, and initial/Dirichlet boundary conditions, $\mathcal{B}_{t,x}$.

Specifically, we use the convection-diffusion equation which is a well-known evolutionary PDE, combines parabolic and hyperbolic equations, is conceptually similar to incompressible Navier-Stokes, and detailed in Appendix A.1.

We aim to learn a surrogate model, i.e, an approximation to the solution, $\hat{u}$:

$$\hat{u} : D_t \times D_x \times \boldsymbol{Q} \mapsto \mathbb{R}^u \tag{1}$$

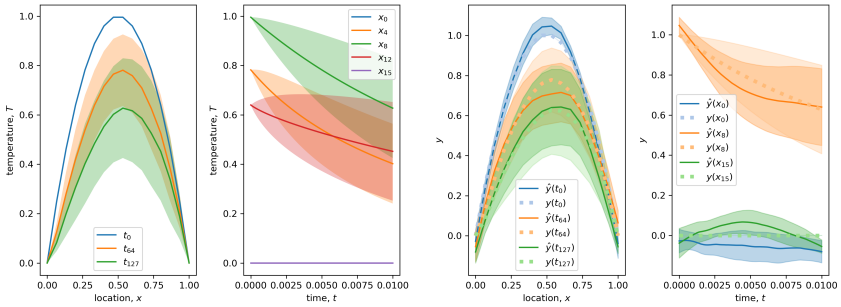

(a) Target solution (b) Spectral PINN solution (ours)

Figure 1: **Spectral PINNs.** Initial results show that the spectral PINNs in Fig. 1b-solid can approximately match the mean (line) and standard deviation (shade) of the target solution in Fig. 1a and Fig. 1b-light-shade-dotted. Importantly the approximated standard deviation also captures the growing trend towards the center location ($x = 0.5$) and growing time, $t = 1$ (far-right-plot, orange).

with the finite set of p mutually independent random parameters, $\boldsymbol{Q}(\omega) = [Q_1(\omega), ..., Q_p(\omega)]$. Each parameter, $Q_i(\omega) : \Omega \mapsto \mathbb{R}$, with realization, $\omega \in \Omega$, has an associated density, $p_{Q_i}(q_i)$, and joint density, $p_Q(q) = \prod_{i=1}^{p} p_{Q_i}(q_i)$. The parameters are in the probability space $(\Omega, \mathcal{F}, P)$ with σ-Algebra, $\mathcal{F}$, sample space, Ω, and probability measure P (see sec. 5.1 in [28]). The solution then maps a distribution of parameters to a distribution of solutions.

2.2 Spectral expansion

One core idea in our paper is to approximate the solution, u, as a spectral expansion, also called polynomial chaos expansion [40, 28]. The spectral expansion projects the solution onto a space of stochastic orthogonal polynomials. The solution then becomes a linear combination of spatiotemporal coefficients and stochastic polynomials. The projection is optimal in the L2 sense [28]. The spectral expansion approximates the solution as,

$$\hat{u}(t, x; \boldsymbol{\xi}) = \sum_{j=0}^{|A|} \hat{C}_{\boldsymbol{\alpha}_j}(t, x) \Psi_{\boldsymbol{\alpha}_j}(\xi_1, ..., \xi_n) \tag{2}$$

with the NN-estimated spectral coefficients, $\hat{C}_{\boldsymbol{\alpha}_j}(t, x) \in \mathbb{R}$. We choose the set of polynomials, $\Psi_A(\boldsymbol{\xi})$, to transform a Gaussian distribution, specifically a multivariate orthogonal Gaussian-Hermite polynomials (as displayed and displayed and detailed in Appendix A.2). The polynomials and spectral coefficients are indexed by the multi-indices, $\boldsymbol{\alpha}_j \in A$, which contain index-pairs into the stochastic and polynomial degree, as detailed in Appendix A.3.

2.3 Spectral PINNs

The NN jointly estimates all $|A|$ spectral coefficients based on 2D spatiotemporal input pairs, $\{t, x\}$,

$$NN_{C_A}(t, x) := \hat{C}_A(t, x) : D_t \times D_x \to \mathbb{R}^{|A|}. \tag{3}$$

Using single input pairs makes our surrogate model mesh-free as the NN can be queried at any location to interpolate the training samples or collocation points. The NN is trained to approximate the spectral coefficients while only using a limited number of measurements as target data. We formulate the mean-squared error (MSE) loss for episode, e, and batch size, B, as:

$$L_e(t_b, x_b) = \frac{1}{B} \sum_{b=1}^{B} ||u(t_b, x_b; \boldsymbol{\xi}_b) - \sum_{j=0}^{|A|} \hat{C}_{\boldsymbol{\alpha}_j}(t_b, x_b) \Psi_{\boldsymbol{\alpha}_j}(\boldsymbol{\xi}_b)||_2^2, \tag{4}$$

where the realizations of Gaussian random vectors, $\boldsymbol{\xi}_b \sim \mathcal{N}(0, 1)^n$ are shared between the target and approximated solution. The batch size is chosen to fit one solution sample, $B = n_t n_x$ with t-grid size, n_t, and x-grid size, n_x. The training and validation/test set are using different realizations of the random vector and the same $\{t, x\}$-pairs as grid points.

3 Preliminary Results

Figure 1 shows that the *Spectral PINNs* in Fig. 1b can successfully approximate the mean and standard deviation of the target solution Fig. 1a. Importantly the explicit formulation as spectral expansion allows us to compute the mean and standard deviation without any sampling as a function of the spectral coefficients, e.g., $\mu_u(t,x) = C_{[0,0,0]}(t,u)$ [28]. We can note that the *Spectral PINNs*-approximated standard deviation captures the growing trend towards the center location ($x = 0.5$) and increasing time ($t = 1$). While the accuracy of estimating the mean and std. dev. could be improved towards regions of high uncertainty ($x = 0.5, t = 0.05$) the plots indicate promise and we expect hyperparameter tuning to further increase the accuracy. Figure 4 shows the approximated spectral coefficients.

We used a 6-layer 2048-unit fully-connected neural network with ResNet blocks, in-/output normalization, and ReLu activation. The network was trained with the ADAM optimizer with learning rate, $lr = 0.001$, and $\beta = [0.9, 0.999]$ for $E = 8$ epochs in ≈1hr on one 5-year old Intel $i7{-}7500U$@2.70GHz quad-core CPU. The target data was generated with $n_t = 128$ temporal and $n_z = 16$ grid points and $\{n_{s,\text{train}} = 800, n_{s,\text{val}} = 200\}$ samples of the solved PDE. The maximum polynomial degree was chosen to be, $n = 2$, s.t. the number of spectral coefficients is $|A| = 10$.

3.1 Linear vs. cubic runtime complexity

Introducing NNs into uncertainty propagation (UP) methods aims to reduce runtime complexity while maintaining accuracy in conducting statistic analysis of high-dimensional PDEs.

Runtime complexity. Classical UP methods can scale cubically with the number of spectral coefficients in the worst-case [28]. For example, stochastic Galerkin requires to solve one system of $J(K+1)$ equations [32] and collocation methods require to solve M systems of size J [34]. Here, J is the number of spectral coefficients, K the number of stochastic polynomials, M the number of training samples and in our case $J = K = |A| = 10,\ M = 1000$. In comparison, inference of our FCNN scales linearly with the number of outputs/spectral coefficients, J, assuming the network architecture is independent of J [44]. In practice, the architecture depends on the number of outputs and future work will investigate this relationship. Further, a full comparison will include a discussion on how training time scales with the number of outputs and experimental plots.

Convergence rate. The final accuracy of collocation methods scales linearly in the number of samples, M [34]. However, the convergence rate towards an accurate solution has poor scaling with the number of stochastic parameters p as $O(M^{-\alpha_{\text{regularity}}/p})$. In comparison, MC sampling converges with $O(M^{-1/2})$ which is faster for a high number of stochastic parameters, p [28]. Future work, will include a discussion on how the accuracy of NNs converges with the number of stochastic training samples, M, and stochastic dimension, p [45].

Memory complexity. In some cases, leveraging NN-based surrogate models can not only reduce computational complexity but also memory complexity [22]. Our network, however, contains $n_w = 2n_{units} + n_{layers}n_{units}^2 + n_{units}|A| = 25M$ weights which occupy as floats $n_{weights}4B \approx 96MB$. As this is quite large, future work, will investigate pruning [46] and mixed-precision [47] methods.

4 Discussion and Future works

We contributed *Spectral PINNs*, a spectral method for fast uncertainty propagation of high-dimensional stochastic PDEs. We demonstrated the first application of NN-based spectral expansions on the convection-diffusion equation. There exist various avenues for future work. First, we would like to extend hyperparameter search to increase accuracy and decrease memory consumption. Second, our work assumes that the stochastic parameters are mutually independent and Gaussian which can be alleviated by an extension to arbitrary polynomial chaos expansion for arbitrary distributions [48]. Third, future work will investigate efficient sampling strategies to reduce the amount of training samples, such as Latin hypercube sampling [49] or Quasi Monte Carlo sampling [50]. Fourth, our work relies purely on FCNNs, but the approximation of as spatiotemporal coefficients naturally lends itself to RNNs [51], CNNs [52], GNNs [53], Transformers [54], or more to exploit spatiotemporal features and propagate uncertaintes of higher-dimensional systems. Lastly, we will continue the analysis of convergence guarantees based on recent work in converging PINNs [45].

References

[1] P. Bauer, A. Thorpe, and G. Brunet, "The quiet revolution of numerical weather prediction," pp. 47–55, sep 2015.

[2] O. Fuhrer, T. Chadha, T. Hoefler, G. Kwasniewski, X. Lapillonne, D. Leutwyler, D. Lüthi, C. Osuna, C. Schär, T. C. Schulthess, and H. Vogt, "Near-global climate simulation at 1 km resolution: establishing a performance baseline on 4888 gpus with cosmo 5.0," *Geosci. Model Dev.*, vol. 11, pp. 1665 – 1681, 2018.

[3] J. P. Boyd, *Chebyshev and Fourier Spectral Methods*, 2nd ed., ser. Dover Books on Mathematics. Mineola, NY: Dover Publications, 2001.

[4] P. Gentine, M. Pritchard, S. Rasp, G. Reinaudi, and G. Yacalis, "Could Machine Learning Break the Convection Parameterization Deadlock?" *Geophysical Research Letters*, vol. 45, no. 11, pp. 5742–5751, jun 2018.

[5] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nature Reviews Physics 2021 3:6*, vol. 3, no. 6, pp. 422–440, may 2021.

[6] M. Raissi, P. Perdikaris, and G. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational Physics*, vol. 378, pp. 686 – 707, 2019.

[7] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, http://www.deeplearningbook.org.

[8] O. L. Maitre and O. M. Knio, *Spectral Methods for Uncertainty Quantification*, 1st ed. Springer Netherlands, 2010, scientific Computation.

[9] M. Reichstein, G. Camps-Valls, B. Stevens, M. Jung, J. Denzler, N. Carvalhais, and Prabhat, "Deep learning and process understanding for data-driven earth system science," *Nature*, vol. 566, pp. 195 – 204, 2019.

[10] S. R. Cachay, E. Erickson, A. F. C. Bucker, E. Pokropek, W. Potosnak, S. Bire, S. Osei, and B. Lütjens, "The world as a graph: Improving el niño forecasts with graph neural networks," 2021.

[11] P. Jiang, N. Meinert, H. J. ao, C. Weisser, S. Holgate, A. Lavin, B. Lütjens, D. Newman, H. Wainwright, C. Walker, and P. Barnard, "Digital twin earth - coasts: Learning a fast and physics-informed surrogate model for coastal floods via neural operators," *arxiv*, 2021.

[12] B. Lusch, J. Kutz, and S. Brunton, "Deep learning for universal linear embeddings of nonlinear dynamics," *Nat. Commun.*, vol. 9, 2018.

[13] S. Greydanus, M. Dzamba, and J. Yosinski, "Hamiltonian neural networks," in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d Alché-Buc, E. Fox, and R. Garnett, Eds. Curran Associates, Inc., 2019, pp. 15 379–15 389.

[14] D. Bau, J.-Y. Zhu, H. Strobelt, A. Lapedriza, B. Zhou, and A. Torralba, "Understanding the role of individual units in a deep neural network," *Proceedings of the National Academy of Sciences*, 2020.

[15] A. T. Mohan, N. Lubbers, D. Livescu, and M. Chertkov, "Embedding hard physical constraints in neural network coarse-graining of 3d turbulence," *ICLR Workshop on AI for Earth Sciences*, 2020.

[16] T. Beucler, M. Pritchard, S. Rasp, J. Ott, P. Baldi, and P. Gentine, "Enforcing analytic constraints in neural networks emulating physical systems," *Phys. Rev. Lett.*, vol. 126, p. 098302, Mar 2021.

[17] B. Lütjens*, B. Leshchinskiy*, C. Requena-Mesa*, F. Chishtie*, N. Díaz-Rodríguez*, O. Boulais*, A. Sankaranarayanan*, A. P. na, Y. Gal, C. Raïssi, A. Lavin, and D. Newman, "Physically-consistent generative adversarial networks for coastal flood visualization," *arXiv e-prints*, 2021, * equal contribution. [Online]. Available: https://arxiv.org/abs/2104.04785

[18] T. Lesort, M. Seurin, X. Li, N. Díaz-Rodríguez, and D. Filliat, "Deep unsupervised state representation learning with robotic priors: a robustness analysis," in *2019 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2019, pp. 1–8.

[19] L. A. Garcia and A. Shigidi, "Using neural networks for parameter estimation in ground water," *Journal of Hydrology*, vol. 318, no. 1, pp. 215–231, 2006.

[20] T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural ordinary differential equations," in *Advances in Neural Information Processing Systems 31*. Curran Associates, Inc., 2018, pp. 6571–6583.

[21] A. Karpatne, W. Watkins, J. Read, and V. Kumar, "Physics-guided Neural Networks (PGNN): An Application in Lake Temperature Modeling," *arXiv e-prints*, p. arXiv:1710.11431, Oct. 2017.

[22] J. Yuval, P. A. O'Gorman, and C. N. Hill, "Use of neural networks for stable, accurate and physically consistent parameterization of subgrid atmospheric processes with good performance at reduced precision," *Geophysical Research Letter*, vol. 48, p. e2020GL091363, 2021.

[23] M. Raissi, "Deep hidden physics models: Deep learning of nonlinear partial differential equations," *Journal of Machine Learning Research*, vol. 19, no. 25, pp. 1–24, 2018.

[24] Z. Long, Y. Lu, and B. Dong, "Pde-net 2.0: Learning pdes from data with a numeric-symbolic hybrid deep network," *Journal of Computational Physics*, vol. 399, p. 108925, 2019.

[25] C. Rackauckas, Y. Ma, J. Martensen, C. Warner, K. Zubov, R. Supekar, D. Skinner, and A. Ramadhan, "Universal differential equations for scientific machine learning," *ArXiv*, vol. abs/2001.04385, 2020.

[26] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané, "Concrete problems in ai safety," 2016.

[27] T. N. Palmer, "Stochastic weather and climate models," pp. 463–471, jul 2019.

[28] R. C. Smith, "Uncertainty quantification: Theory, implementation, and applications," in *Computational science and engineering*. SIAM, 2013, p. 382.

[29] P. Kidger, J. Foster, X. Li, H. Oberhauser, and T. Lyons, "Neural SDEs as Infinite-Dimensional GANs," *International Conference on Machine Learning*, 2021.

[30] X. Li, T.-K. L. Wong, R. T. Q. Chen, and D. Duvenaud, "Scalable Gradients for Stochastic Differential Equations," 2020.

[31] L. Yang, X. Meng, and G. Karniadakis, "B-pinns: Bayesian physics-informed neural networks for forward and inverse pde problems with noisy data," *J. Comput. Phys.*, vol. 425, p. 109913, 2021.

[32] I. Babuška, R. Tempone, and G. E. Zouraris, "Galerkin finite element approximations of stochastic elliptic partial differential equations," *SIAM Journal on Numerical Analysis*, vol. 42, no. 2, pp. 800–825, 2005.

[33] W. K. Hastings, "Monte Carlo sampling methods using Markov chains and their applications," *Biometrika*, vol. 57, no. 1, pp. 97–109, 04 1970.

[34] I. Babuška, F. Nobile, and R. Tempone, "A stochastic collocation method for elliptic partial differential equations with random input data," *SIAM J. Numer. Anal.*, vol. 45, no. 3, p. 1005–1034, May 2007.

[35] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems*, vol. 27. Curran Associates, Inc., 2014, pp. 2672–2680.

[36] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," in *2nd International Conference on Learning Representations, ICLR*, 2014.

[37] J. Sohl-Dickstein, E. A. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *Proceedings of the 32nd International Conference on Machine Learning*, vol. 37. JMLR.org, 2015, pp. 2256–2265.

[38] D. Rezende and S. Mohamed, "Variational inference with normalizing flows," in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, F. Bach and D. Blei, Eds., vol. 37. Lille, France: PMLR, 07–09 Jul 2015, pp. 1530–1538.

[39] N. Wiener, "The homogeneous chaos," *American Journal of Mathematics*, vol. 60, no. 4, pp. 897–936, 1938.

[40] T. Crestaux, O. Le Maı^tre, and J.-M. Martinez, "Polynomial chaos expansion for sensitivity analysis," *Reliability Engineering and System Safety*, vol. 94, no. 7, pp. 1161–1172, 2009, special Issue on Sensitivity Analysis.

[41] B. Sudret, "Global sensitivity analysis using polynomial chaos expansions," *Reliability Engineering and System Safety*, vol. 93, no. 7, pp. 964–979, 2008, bayesian Networks in Dependability.

[42] S. Oladyshkin and W. Nowak, "Data-driven uncertainty quantification using the arbitrary polynomial chaos expansion," *Reliability Engineering and System Safety*, vol. 106, pp. 179–190, 2012.

[43] D. Zhang, L. Lu, L. Guo, and G. E. Karniadakis, "Quantifying total uncertainty in physics-informed neural networks for solving forward and inverse stochastic problems," *Journal of Computational Physics*, vol. 397, p. 108850, 2019.

[44] K. Fredenslund, "Computational complexity of neural networks," last accessed 09/2021. [Online]. Available: https://kasperfred.com/series/introduction-to-neural-networks/computational-complexity-of-neural-networks

[45] S. Mishra and R. Molinaro, "Estimates on the generalization error of physics informed neural networks (pinns) for approximating pdes," 2021.

[46] L. Liebenwein, C. Baykal, H. Lang, D. Feldman, and D. Rus, "Provable filter pruning for efficient neural networks," in *International Conference on Learning Representations*, 2020.

[47] M. Courbariaux, Y. Bengio, and J.-P. David, "Training deep neural networks with low precision multiplications," 2015.

[48] X. Wan and G. E. Karniadakis, "Multi-element generalized polynomial chaos for arbitrary probability measures," *SIAM J. Sci. Comput.*, vol. 28, no. 3, p. 901–928, Mar. 2006.

[49] M. D. McKay, R. J. Beckman, and W. J. Conover, "A comparison of three methods for selecting values of input variables in the analysis of output from a computer code," *Technometrics*, vol. 21, no. 2, pp. 239–245, 1979.

[50] W. J. Morokoff and R. E. Caflisch, "Quasi-monte carlo integration," *Journal of Computational Physics*, vol. 122, no. 2, pp. 218–230, 1995.

[51] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, *Learning Internal Representations by Error Propagation*. Cambridge, MA, USA: MIT Press, 1986, p. 318–362.

[52] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, Eds., vol. 25. Curran Associates, Inc., 2012.

[53] W. L. Hamilton, "Graph representation learning," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 14, no. 3, pp. 1–159, 2020.

[54] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," *Association for Computational Linguistics Human Language Technologies (NAACL-HLT)*, vol. 1, pp. 4171–4186, 2019.

A Appendix

A.1 The local convection-diffusion equation

We are given the local convection-diffusion equation which models the temperature distribution in a vertical ocean column over time,

$$f = \frac{\delta T(t,z;\omega)}{\delta t} + w\frac{\delta T(t,z;\omega)}{\delta z} - \frac{\delta}{\delta z}\left(\kappa(z;\omega)\frac{\delta T(t,z;\omega)}{\delta z}\right) \tag{5}$$

with height, $z \in D_z = [0,1]$, time, $t \in D_t = [0, 0.01]$, source, $f = 0$, noise, $\omega \in \Omega$, temperature, $T : D_t, D_z \to \mathbb{R}$, stochastic diffusivity, $Q = [\kappa(z,\omega)]$, constant vertical velocity, $w = 100$, initial temperature, $T(0,z) = -4(z-0.5)^2 + 1$, and periodic Dirichlet boundary conditions, $T(t,0) = T(t,1) = 0$.

We assume that the distribution over the diffusivity is known, for example, through domain knowledge, data assimilation, or Bayesian parameter estimation. Specifically, the diffusivity is assumed to follow

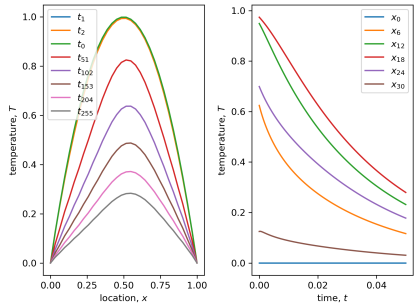

Figure 2: **Local convection diffusion equation.** The local convection diffusion equation simulates how the initial temperature profile, $T(t = 0, z)$ (green), is distributed over time. We can observe that the positive diffusivity, $\kappa > 0$, flattens the temperature curve (left) and the vertical velocity, $w > 0$, shifts the curve to the right over time.

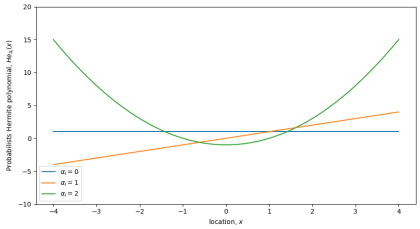

Figure 3: **Hermite Polynomials.** The basis functions for the first three Hermite polynomials, $\text{He}_i(x)$ for $i = \{0, 1, 2\}$. As the

an exponential Gaussian process (GP) with $\kappa(z;\omega) = \exp(Y_\kappa(z;\omega))$. The GP, $Y_\kappa(z;\omega)$, is defined by mean, $\mu_{Y_\kappa} = 1000$, correlation length, $L = 0.3$, variance, $\sigma_{Y_\kappa} = 1.0$, exponent, $p_{\text{GP}} = 1.0$, and a covariance kernel that is similar to the non-smooth Ornstein-Uhlenbeck kernel:

$$\text{Cov}_{Y_\kappa}(z_1, z_2) = \sigma^2_{Y_\kappa} \exp(-\frac{1}{p_{GP}}(\frac{|z_1 - z_2|}{L})^p). \tag{6}$$

We generated $n = 1000$ samples of the stochastic convection-diffusion equation via a 2nd-order finite difference scheme in Python.

A.2 Hermite polynomials

We choose the set of polynomials, $\Psi_A(\boldsymbol{\xi})$, to be a set of multivariate orthogonal Gaussian-Hermite polynomials, as displayed in Fig. 3:

$$\begin{aligned} \Psi_{\boldsymbol{\alpha}_j}(\xi_1, ..., \xi_n) &= \Pi_{i=1}^n \psi_{\alpha_{ji}}(\xi_i), \\ &= \Pi_{i=1}^n \text{He}_{\alpha_{ji}}(\xi_i), \\ \text{with } \xi_i &\sim \mathcal{N}(0, 1), \end{aligned} \tag{7}$$

with the one-term (monic) polynomials, $\psi_{\alpha_{ji}}$ of polynomial degree, α_{ji}, and maximum degree, n. We are choosing the random vector of each stochastic dimension, $i \in \{0, ..., n\}$, to be a Gaussian, $\xi_i \sim \mathcal{N}(0, 1)$ and use the associated probabilists' Hermite polynomials, $\text{He}_{\alpha_{ji}}$.

A.3 Multi-indices

The polynomial degrees are given by the total-degree multi-index set, $A = \{\boldsymbol{\alpha}_j \in \mathbb{N}_0^n : ||\boldsymbol{\alpha}_j||_1 = \sum_{i=1}^n \alpha_{ji} \leq n\}$. For example, $A = \{[0, 0, 0], [0, 0, 1], [0, 1, 0], [1, 0, 0], [0, 0, 2], [0, 1, 1], [0, 2, 0], [1, 0, 1], [1, 1, 0], [2, 0, 0]\}$ for $n = 2$. The vector of polynomial degrees or multi-index is given by $\boldsymbol{\alpha}_j \in A$ with $j \in \{0, ..., |A|\}$ [28].

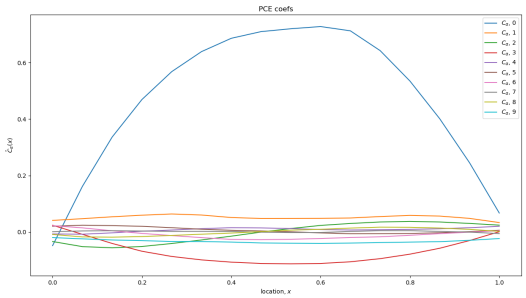

Figure 4: **Learned spectral coefficients.** The neural network outputs the first 10 spectral coefficients. We can observe that the first coefficient, $C_{[0,0,0]}$ approximates the mean and is of significantly larger scale than the other coefficients.

A.4 Learned spectral coefficients

Figure 4 shows the approximated spectral coefficients.