
Scalable Deep-Learning-Accelerated Topology Optimization for Additively Manufactured Materials

Sirui Bi

Computational Sciences and Engineering Division
Oak Ridge National Laboratory, Oak Ridge, TN 37830
bis1@ornl.gov

Jiixin Zhang

Computer Science and Mathematics Division
Oak Ridge National Laboratory, Oak Ridge, TN 37830
zhangj@ornl.gov

Guannan Zhang

Computer Science and Mathematics Division
Oak Ridge National Laboratory, Oak Ridge, TN 37830
zhangg@ornl.gov

Abstract

Topology optimization (TO) is a popular and powerful computational approach for designing novel structures, materials, and devices. Two computational challenges have limited the applicability of TO to a variety of industrial applications. First, a TO problem often involves a large number of design variables to guarantee sufficient expressive power. Second, many TO problems require a large number of expensive physical model simulations, and those simulations cannot be parallelized. To address these issues, we propose a general scalable deep-learning (DL) based TO framework, referred to as SDL-TO, which utilizes parallel schemes in high performance computing (HPC) to accelerate the TO process for designing additively manufactured (AM) materials. Unlike the existing studies of DL for TO, our framework accelerates TO by learning the iterative history data and simultaneously training on the mapping between the given design and its gradient. The surrogate gradient is learned by utilizing parallel computing on multiple CPUs incorporated with a distributed DL training on multiple GPUs. The learned TO gradient enables a fast online update scheme instead of an expensive update based on the physical simulator or solver. Using a local sampling strategy, we achieve to reduce the intrinsic high dimensionality of the design space and improve the training accuracy and the scalability of the SDL-TO framework. The method is demonstrated by benchmark examples and AM materials design for heat conduction. The proposed SDL-TO framework shows competitive performance compared to the baseline methods but significantly reduces the computational cost by a speed up of around 8.6x over the standard TO implementation.

1 Introduction

Topology optimization (TO) has been extensively applied to solving complex engineering design problems with a wide range of industrial applications, e.g. aerospace, building, and automotive. Unfortunately, TO has a major challenge in computational expense because it is an iterative and

sequential procedure, which is not straightforward to be palatalized. A typical large-scale TO problem may involve hundreds or even thousands of design iterations, and for each iteration, the physical response needs to be solved to compute the gradient information through sensitivity analysis. To handle large-scale TO problem, which involves millions of design variables [1, 3, 15], the associated computational cost is highly prohibitive and will be larger and larger in the future since the problem are becoming more complex with more details [1]. In general, there are several types of methods to improve the computational efficiency, including algorithm improvements [4, 20], e.g., advanced iterative solvers [21, 16], parallel computing [1, 2] including GPU computing [22, 44, 39], and deep learning techniques [32, 34], e.g., reparameterization [18, 17].

Another typical strategy to improve the efficiency of TO simulations is to build surrogate modeling, which has been developed in a wide variety of contexts and disciplines, see [6, 9, 10, 12] and the reference therein. Surrogate modeling practice seeks to approximate the response of an original function (the loss function in this work), which is typically computationally expensive, by a cheaper-to-run surrogate. The loss function can then be evaluated by evaluating the surrogate directly without running expensive forward model executions. Compared to conventional TO algorithms, this approach is advantageous that it significantly reduces the number of forward model executions at a desired accuracy and allows evaluating the loss function in parallel. Several methods can be employed to construct the surrogate systems, including least-squares projection [26, 25, 11, 24], sparse interpolation [29, 28, 14], and compressive sensing [13, 23, 40, 42, 33, 31], etc. However, unlike the standard surrogate modeling task, when TO does not require an accurate surrogate model in the entire parameter space, because an optimizer only goes down a single path from the initial state to the final state to find the optimum. Thus, it is wasteful to build a global surrogate to the loss function. On the other hand, as TO is often a high-dimensional optimization problem, it is also challenging to build a global surrogate with sufficient accuracy even with the state-of-the-art neural network models.

In this work, we first propose a general scalable framework integrating deep learning and parallel computing to accelerate the TO process for designing additively manufactured materials. The key idea is to dig into the iterative history data and use deep neural networks (DNN) to learn a surrogate gradient given a specific design. The learning step is achieved by utilizing parallel computing on multiple CPUs and a distributed DNN training on multiple GPUs, incorporated with a novel local sampling strategy, which aims to reduce the dimensionality to capture a low-dimensional search space. The proposed SDL-TO framework is demonstrated by several benchmark examples and compared with the start-of-the-art baseline methods. In the large-scale AM materials design example, SDL-TO significantly saves the computational cost and achieves a speed up of around 8.6x over the standard TO implementation.

2 SDL-TO: Scalable Deep Learning Accelerated Topology Optimization

In this paper, we propose a scalable deep learning framework to accelerate TO, but we emphasize that it can be generally applied to a wide range of computational design optimization problems, specifically large-scale, high-dimensional, computationally intensive problems in computational sciences and engineering (CSE). The core idea is using an accurately learned gradient by DNN instead of the true gradient that is iteratively evaluated at each step to accelerate the optimization process. This is achieved by utilizing a local sampling strategy incorporating a scalable implementation on multiple CPUs and GPUs. Figure 1 gives an overview of this process. Note that the proposed framework is universal in the sense that it can work with any gradient-based optimization based on computationally intensive physical models or simulators.

Topology optimization formulation TO is a mathematical method that aims to optimize materials layout defined on a design domain with given boundary conditions, loads and volume constraints to minimize structural compliance C (or equivalently the least strain energy). In this work, we use the modified Solid Isotropic Material with Penalization (SIMP) approach [36] with density-based approach to TO, where each element e is assigned a density x_e that determines its Young’s modulus:

$$E_e(x_e) = E_{\min} + x_e^p(E_0 - E_{\min}), \quad x_e \in [0, 1] \quad (1)$$

where E_0 is the stiffness of the material, E_{\min} is a very small stiffness assigned to void regions to prevent the stiffness matrix becoming singular. The modified SIMP approach differs from the classical SIMP approach [7], where elements with zero stiffness are avoided by using a small value.

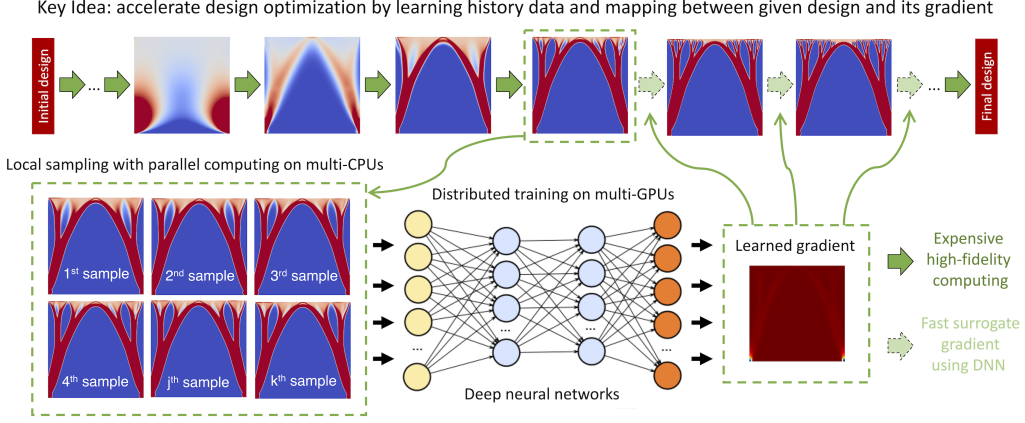


Figure 1: SDL-TO framework to accelerating TO by learning the iterative history data and mapping between given design and its gradient. Using local sampling strategy with distributed computing and training on multiple CPUs and GPUs, the learned surrogate gradient enables an online update instead of expensive computing using a high-fidelity simulator or solver.

The modified mathematical formulation of the TO problem is

$$\min_{\mathbf{x}} : C(\mathbf{x}) = \mathbf{U}^T \mathbf{K} \mathbf{U} = \sum_{e=1}^N E_e(x_e) \mathbf{u}_e^T \mathbf{k}_e \mathbf{u}_e, \quad \text{s.t.} \quad \mathbf{K} \mathbf{U} = \mathbf{F}, V(\mathbf{x}) = V_0, 0 \leq \mathbf{x} \leq 1 \quad (2)$$

where \mathbf{x} is the vector of design variables, C is the structural compliance, \mathbf{K} is the global stiffness matrix, \mathbf{U} and \mathbf{F} are the global displacement and force vectors respectively, \mathbf{u}_e and \mathbf{k}_e are the element displacement vector and stiffness matrix respectively, N is the number of elements used to discretize the design domain Ω , $V(\mathbf{x})$ and V_0 are the material volume and design domain volume respectively, ζ is the prescribed volume fraction, and p is the penalization power coefficient (typically $p = 3$).

TO baseline method The computationally limiting step in standard TO is the gradient estimation using an adjoint method [8], which typically needs to solve $\mathbf{U} = \mathbf{K}^{-1} \mathbf{F}$ using the finite element method (FEM). The gradient is inserted into the Method of Moving Asymptotes (MMA) [38] algorithm that is the state-of-the-art optimizer, which has been extensively demonstrated to be versatile and well suited for a wide range TO problems. We implemented this algorithm using NumPy and compare our method with the standard TO process using MMA in a single CPU. Note that our scalable framework aims to accelerate the iterative optimization process and the comparison is based on the same solver and TO algorithm. The improvement of the TO algorithm self, the linear solver, or the physical simulator using parallel computing is out of the scope of this study.

Online update with a local sampling strategy To large-scale TO problems, typically involving high resolution, the design space is so huge that it is very challenging to directly map the given design and its gradient using random samples. It is a non-trivial task to use a global surrogate model to cover the entire high-dimensional space in this case. To address this issue, we propose a local sampling strategy that allows us to capture the optimal low-dimensional gradient direction by generating random samples around the current design with a few history steps before. Since the standard TO algorithm uses gradient descent for a local search, the adjacent design does not show significant changes except for the initial stage. Therefore, we propose a local sampling strategy using local samples that are drawn from a Gaussian distribution

$$\tilde{\mathbf{x}}_{rs}^{j+1} \sim \mathcal{N}(\mathbf{x}_{\min}^{j+1} + (\mathbf{x}_{\max}^{j+1} - \mathbf{x}_{\min}^{j+1})/2, \sigma) \quad (3)$$

where σ is the standard deviation which is used to control the radius of the local search domain, and the upper and lower bound are defined by

$$\tilde{\mathbf{x}}_{\max}^{j+1} = \max \left\{ \tilde{\mathbf{x}}^{(j)}, \tilde{\mathbf{x}}^{(j-1)}, \dots, \tilde{\mathbf{x}}^{(j-w)} \right\}, \quad \tilde{\mathbf{x}}_{\min}^{j+1} = \min \left\{ \tilde{\mathbf{x}}^{(j)}, \tilde{\mathbf{x}}^{(j-1)}, \dots, \tilde{\mathbf{x}}^{(j-w)} \right\} \quad (4)$$

where w is a lookback window size, which means how many history steps we will look back from the current state $\tilde{\mathbf{x}}^{j+1}$. As shown in Figure 1, we have a number of designs and their corresponding

gradients from these local samples. A deep neural network is trained to map the design and its gradient and we therefore learn a accurate gradient for any specific design within this local search domain. Once we have the learned gradient in hand, we may continue gradient-based optimization without any expensive simulations, e.g., FEM, but a critical question is bought up: *how many steps can we move forward to the optimal design?* This is determined by defining a threshold λ^* of the cosine distance d^c between current design $\tilde{\mathbf{x}}^{j+1}$ and future design $\tilde{\mathbf{x}}^{j+m}$. If $d^c > \lambda^*$, the *learning* step consisting of simulation and training will be called, otherwise, the learned surrogate gradient will be used for online next step update of the design.

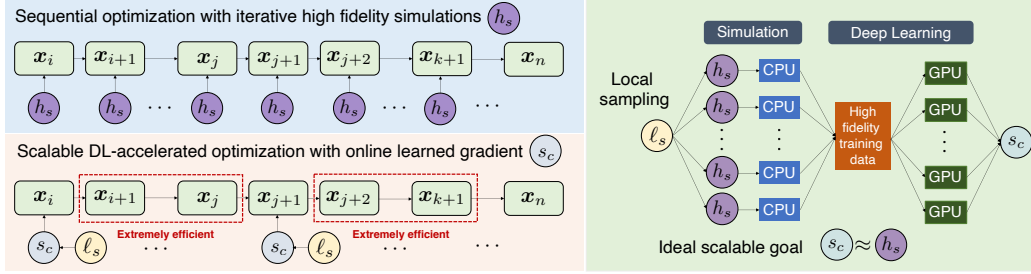


Figure 2: Scalable implementation of the SDL-TO framework on multiple CPUs and GPUs

Scalable implementation of the SDL-TO framework A core contribution of the SDL-TO framework is to utilize high performance computing (HPC) sources (multiple CPUs and GPUs) to accelerate the sequentially iterative TO process in a parallel scheme. Figure 2 shows the schema of our scalable implementation. Compared with the sequential TO iterative process $\mathbf{x}_i, \mathbf{x}_{i+1}, \dots, \mathbf{x}_n$ with high fidelity FEM simulation h_s at each step, we accelerate optimization with an online update using the learned gradient s_c incorporating with a local sampling strategy ℓ_s , where local samples of h_s are generated from a Gaussian distribution in Eq. (3) and sent to multiple CPUs for parallel simulations. Then we collect all h_s data together for the distributed deep learning training on multiple GPUs to obtain a fast and accurate mapping between the current design and its gradient. In such a way, we only need to call a very small number of *learning* steps, but use the learned gradient s_c for a fast design update. For example, we call the learning step at \mathbf{x}_{j+1} and then we can use the learned gradient for multiple later updating steps $\mathbf{x}_{j+2}, \dots, \mathbf{x}_{k+1}$ ($k > j + 1$). Ideally, the total computational cost of s_c equals h_s but it needs to be fine-tuned and optimized due to the communication cost and a requirement of large enough computing resources. In this work, we use `mpi4py` for implementing the parallel simulation on multiple GPUs and `Horovod` [35] for the distributed training on multiple GPUs. All the parallel scheme is conducted and tested on the ORNL OLCF Summit supercomputer, where each Summit node consists of 2 IBM Power9 CPUs, 6 NVIDIA V100 GPUs, and NVLink for CPU-CPU and CPU-GPU communications.

3 Examples

Benchmark bridge structure We first demonstrate our proposed framework on a benchmark bridge design problem. As shown in Figure 3, there are 80% void materials with $E_1 = 0.001$ and 20% solid materials with $E_2 = 1.0$ so that our goal is to optimize the layout of the materials to achieve minimum structural compliance of the support structure given a unit uniform load on the top. The SDL-TO framework shows the competitive performance of the objective loss values on two resolutions: 480×160 pixels (mesh grids) and 960×320 pixels (mesh grids). When the resolution is higher, our method outperforms the sequential TO (Seq-TO) method and shows smaller structures that benefit to improve the overall structural stiffness (the smaller objective (Obj) value, the better performance).

The performance can be also demonstrated by a quantitative comparison of SDL-TO and Seq-TO in terms of iteration history and time-to-solution metric, as shown in Figure 4. The objective (loss) values are very close in terms of the iteration process. To illustrate the small differences, Figure 4(b) and (c) display the zoom-in results of two resolutions in the 75-200 iteration steps. Figure 4(d) shows the scalability performance in terms of time-to-solution. The Seq-TO is compared with multiple numbers of nodes including 16, 32, 64, and 128. It is interesting to note that the SDL-TO framework




Figure 3: Comparing baselines on the bridge structure design example on two mesh grids (480×160 pixels and 960×320 pixels). We found that the SDL-TO shows competitive performance on both the objective loss and the final design compared to the sequential TO that uses the MMA optimizer.

shows stronger scalability in the high-resolution case than the low-resolution case. This is probably due to a higher ratio of the communication overhead in the low-resolution case.




Figure 4: Quantitative comparison of the SDL-TO and the Seq-TO on objective loss and time-to-solution. (a) Iterative history of two mesh grids, zoom-in results of (b) 480×160 pixels and (c) 960×320 pixels, and (d) time-to-solution performance as the number of nodes increases.

Two materials heat conduction This is a large-scale TO problem that aims to address a challenging problem in additive manufacturing (AM): how to effectively distribute heating by optimizing the layout of two materials consisting of 40% good conductor with $k_1 = 1.0$ and 60% bad conductor with $k_2 = 0.001$, as shown in Figure 5. Considering the AM constraints in practice, we use a very high resolution, 1024×1024 mesh grids in this case for a large-scale demonstration. Thus, the conventional sequential TO process is more computationally intensive than the bridge design example discussed above. Here, we assume to use a total of 200 iterations for the optimization design and compare the performance of the final design and time-to-solution between the Seq-TO method and SDL-TO method.

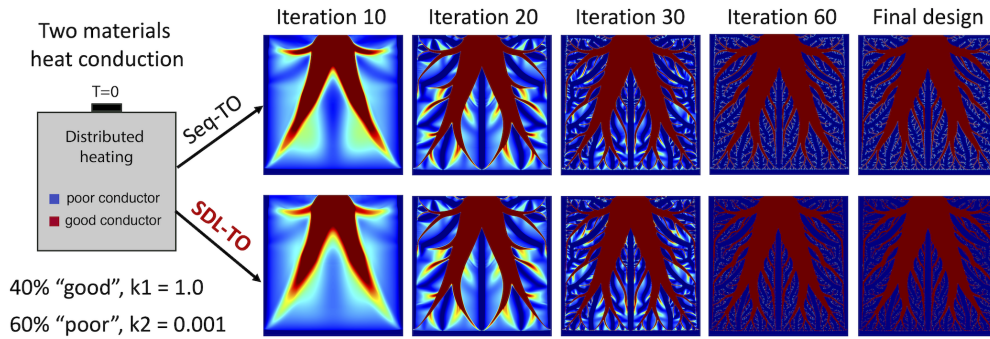


Figure 5: Two materials heat conduction design problem. Compared with the Seq-TO algorithm that uses 200 iterations, the SDL-TO demonstrates a very close performance but only using 12 gradient learning steps. The other updates are achieved by using the learned surrogate gradient.

By comparison of the different iterative designs, specifically on iteration 10, 20, 30, and 60, as shown in Figure 5, we can demonstrate that the SDL-TO framework using surrogate gradients learned by deep learning presents competitive performance on the final objective. For the iterative objective values shown in Figure 6, there are also almost no differences between the sequential TO methods and the proposed scalable TO approach. However, the SDL-TO framework only uses 12 gradient learning steps and then uses the surrogate gradient for a fast design update. In terms of the time-to-solution performance, the SDL-TO method outperforms the sequential TO method. As shown in Figure 6, the SDL-TO framework shows a speed up of 1.08x, 2.23x, 4.58x, and 8.60x using 16, 32, 64, and 128 nodes respectively.

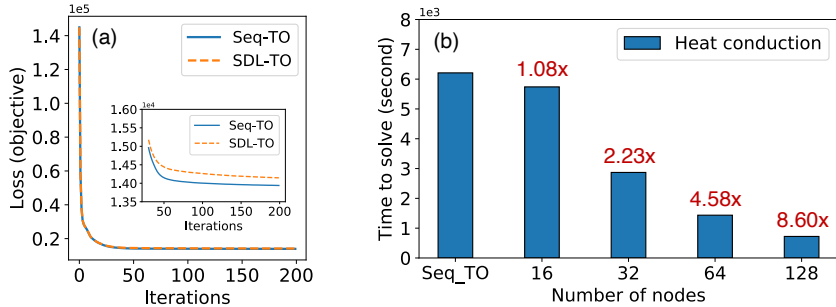


Figure 6: Iteration history and scalability performance. The SDL-TO framework demonstrates a speed up of 8.6x accelerating over the standard TO implementation.

4 Related work

Parameterization TO Hoyer et al. [17] proposed a neural reparameterization that improves topology optimization. They considered the use of implicit bias over functions by neural networks to improve the parameterization of TO and optimize the parameters of the neural networks instead of directly optimizing the densities on the mesh grid. This method can potentially achieve a better design with smaller objective loss values but it did not show how to accelerate the TO process within the proposed framework. Kallioras et al. [18] proposed a new representation using deep belief neural (DBN) coupled with the SIMP approach to accelerate the TO process. This method reduces the number of iterations but the overall time-to-solution is not ideal due to the relatively large expense of the DBN training and data generation.

Neural networks and TO Several recent work focus on the use of convolutional neural network (CNN) [5, 45, 37, 41, 19] or generative adversarial network (GAN) [27, 43, 30] to replace the standard TO methods. Most of them start from creating a dataset of structures/materials via the standard TO and then train a machine learning model based on the dataset. These methods can only reproduce their training data that needs a large computational cost in advance. The trained model is strongly limited by the training dataset and is difficult to handle a general problem. In contrast, our method pursues a general framework to accelerate the TO process by using the iterative history data rather than the end of the final design.

5 Conclusion

In this work, we develop a scalable framework that combined deep learning and parallel computing to accelerate the TO process for designing large-scale additively manufactured materials. The novel contribution is to deeply understand the iterative history data and utilize deep neural networks (DNN) to learn an accurate and fast surrogate gradient instead of the true gradient that is often time-consuming. The learning process is achieved by utilizing parallel computing on multiple CPUs and a distributed DNN training on multiple GPUs, incorporated with a local sampling strategy that enables to capture of the low-dimensional search space. The proposed SDL-TO framework is demonstrated by benchmark examples and compared with the standard sequential TO baseline methods. In the large-scale AM materials design example, SDL-TO shows competitive performance but significantly reduces the computational cost by achieving a speed up of 8.6x over the standard TO implementation.

6 Acknowledgments

This work was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research (ASCR), Applied Mathematics program under contract ERKJ352, ERKJ369; and by the Artificial Intelligence Initiative at the Oak Ridge National Laboratory (ORNL). This work used resources of the Oak Ridge Leadership Computing Facility, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

References

- [1] Niels Aage, Erik Andreassen, Boyan S Lazarov, and Ole Sigmund. Giga-voxel computational morphogenesis for structural design. *Nature*, 550(7674):84–86, 2017.
- [2] Niels Aage and Boyan S Lazarov. Parallel framework for topology optimization using the method of moving asymptotes. *Structural and multidisciplinary optimization*, 47(4):493–505, 2013.
- [3] Joe Alexandersen, Ole Sigmund, and Niels Aage. Large scale three-dimensional topology optimisation of heat sinks cooled by natural convection. *International Journal of Heat and Mass Transfer*, 100:876–891, 2016.
- [4] Oded Amir and Ole Sigmund. On reducing computational effort in topology optimization: how far can we go? *Structural and Multidisciplinary Optimization*, 44(1):25–29, 2011.
- [5] Saurabh Banga, Harsh Gehani, Sanket Bhilare, Sagar Patel, and Levent Kara. 3d topology optimization using convolutional neural networks. *arXiv preprint arXiv:1808.07440*, 2018.
- [6] J. Beck, R. Tempone, F. Nobile, and L. Tamellini. On the optimal polynomial approximation of stochastic PDEs by Galerkin and collocation methods. *Math. Models and Methods Appl. Sci.*, 22(1250023), 2012.
- [7] Martin P Bendsøe. Optimal shape design as a material distribution problem. *Structural optimization*, 1(4):193–202, 1989.
- [8] Martin Philip Bendsoe and Ole Sigmund. *Topology optimization: theory, methods, and applications*. Springer Science & Business Media, 2013.
- [9] A. Chkifa, A. Cohen, R. DeVore, and C. Schwab. Sparse adaptive taylor approximation algorithms for parametric and stochastic elliptic PDEs. *Modél. Math. Anal. Numér.*, 47(1):253–280, 2013.
- [10] A. Chkifa, A. Cohen, and C. Schwab. Breaking the curse of dimensionality in sparse polynomial approximation of parametric PDEs. *J. Math. Pures Appl.*, pages 1–25, 2014.
- [11] Abdellah Chkifa, Albert Cohen, Giovanni Migliorati, Fabio Nobile, and Raul Tempone. Discrete least squares polynomial approximation with random evaluations—application to parametric and stochastic elliptic PDEs. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(3):815–837, 2015.
- [12] Albert Cohen, Ronald DeVore, and Christoph Schwab. Convergence Rates of Best N-term Galerkin Approximations for a Class of Elliptic sPDEs. *Foundations of Computational Mathematics*, 10(6):615–646, July 2010.
- [13] A. Doostan and H. Owhadi. A non-adapted sparse approximation of pdes with stochastic inputs. *Journal of Computational Physics*, 230:3015–3034, 2011.
- [14] M. S. Eldred, C. G. Webster, and P. Constantine. Evaluation of non-intrusive approaches for Wiener-Askey generalized polynomial chaos. In *Proceedings of the Structures, Structural Dynamics, and Materials Conference (10th AIAA Non-Deterministic Approaches Conference)*, volume 1892, pages 1–23. AIAA/ASME/ASCE/AHS/ASC, April 2008.
- [15] Anton Evgrafov, Cory J Rupp, Kurt Maute, and Martin L Dunn. Large-scale parallel topology optimization using a dual-primal substructuring solver. *Structural and Multidisciplinary Optimization*, 36(4):329–345, 2008.

- [16] Federico Ferrari and Ole Sigmund. A new generation 99 line matlab code for compliance topology optimization and its extension to 3d. *arXiv preprint arXiv:2005.05436*, 2020.
- [17] Stephan Hoyer, Jascha Sohl-Dickstein, and Sam Greydanus. Neural reparameterization improves structural optimization. *arXiv preprint arXiv:1909.04240*, 2019.
- [18] Nikos Ath Kallioras, Georgios Kazakis, and Nikos D Lagaros. Accelerated topology optimization by means of deep learning. *Structural and multidisciplinary optimization,(under review)*, 2020.
- [19] Hunter T Kollmann, Diab W Abueidda, Seid Koric, Erman Guleryuz, and Nahil A Sobh. Deep learning for topology optimization of 2d metamaterials. *Materials & Design*, 196:109098, 2020.
- [20] Weichen Li, Phanish Suryanarayana, and Glaucio H Paulino. Accelerated fixed-point formulation of topology optimization: Application to compliance minimization problems. *Mechanics Research Communications*, 103:103469, 2020.
- [21] Zhongyuan Liao, Yu Zhang, Yingjun Wang, and Weihua Li. A triple acceleration method for topology optimization. *Structural and Multidisciplinary Optimization*, 60(2):727–744, 2019.
- [22] Jesús Martínez-Frutos and David Herrero-Pérez. Gpu acceleration for evolutionary topology optimization of continuum structures using isosurfaces. *Computers & Structures*, 182:119–136, 2017.
- [23] L. Mathelin and K. Gallivan. A compressed sensing approach for partial differential equations with random input data. *Commun. Comput. Phys.*, 12:919–954, 2012.
- [24] G Migliorati. *Polynomial approximation by means of the random discrete L^2 projection and application to inverse problems for PDEs with stochastic data*. PhD thesis, Ecole Polytechnique, Inria France, POLITECNICO DI MILANO, October 2014.
- [25] G Migliorati, F Nobile, E von Schwerin, and R Tempone. Approximation of Quantities of Interest in Stochastic PDEs by the Random Discrete L^2 Projection on Polynomial Spaces. *SIAM Journal on Scientific Computing*, 35(3):A1440–A1460, January 2013.
- [26] G Migliorati, F Nobile, E von Schwerin, and R Tempone. Analysis of Discrete L^2 Projection on Polynomial Spaces with Random Evaluations. *Foundations of Computational Mathematics*, 14(3):419–456, Jun 2014.
- [27] Zhenguo Nie, Tong Lin, Haoliang Jiang, and Levent Burak Kara. Topologygan: Topology optimization using generative adversarial networks based on physical fields over the initial domain. *arXiv preprint arXiv:2003.04685*, 2020.
- [28] F Nobile, R Tempone, and C G Webster. An anisotropic sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM Journal on Numerical Analysis*, 46(5):2411–2442, 2008.
- [29] F. Nobile, R. Tempone, and C. G. Webster. A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM Journal on Numerical Analysis*, 46(5):2309–2345, 2008.
- [30] Sangeun Oh, Yongsu Jung, Seongsin Kim, Ikjin Lee, and Namwoo Kang. Deep generative design: Integration of topology optimization and generative models. *Journal of Mechanical Design*, 141(11), 2019.
- [31] Ji Peng, Jerrad Hampton, and Alireza Doostan. A weighted ℓ_1 -minimization approach for sparse polynomial chaos expansions. *Journal of Computational Physics*, 267:92–111, 2014.
- [32] Chao Qian and Wenjing Ye. Accelerating gradient-based topology optimization design with dual-model neural networks. *arXiv preprint arXiv:2009.06245*, 2020.
- [33] H. Rauhut and C. Schwab. Compressive sensing Petrov-Galerkin approximation of high dimensional parametric operator equations. *arXiv*, 2014.

- [34] Hidenori Sasaki and Hajime Igarashi. Topology optimization accelerated by deep learning. *IEEE Transactions on Magnetics*, 55(6):1–5, 2019.
- [35] Alexander Sergeev and Mike Del Balso. Horovod: fast and easy distributed deep learning in tensorflow. *arXiv preprint arXiv:1802.05799*, 2018.
- [36] Ole Sigmund. Morphology-based black and white filters for topology optimization. *Structural and Multidisciplinary Optimization*, 33(4-5):401–424, 2007.
- [37] Ivan Sosnovik and Ivan Oseledets. Neural networks for topology optimization. *Russian Journal of Numerical Analysis and Mathematical Modelling*, 34(4):215–223, 2019.
- [38] Krister Svanberg. The method of moving asymptotes—a new method for structural optimization. *International journal for numerical methods in engineering*, 24(2):359–373, 1987.
- [39] Zhaohui Xia, Yingjun Wang, Qifu Wang, and Chao Mei. Gpu parallel strategy for parameterized lsm-based topology optimization using isogeometric analysis. *Structural and Multidisciplinary Optimization*, 56(2):413–434, 2017.
- [40] L. Yan, L. Guo, and D. Xiu. Stochastic collocation algorithms using ℓ_1 -minimization. *International Journal for Uncertainty Quantification*, 2(3):279–293, 2012.
- [41] Liang Yang, Zesheng Kang, Xiaochun Cao, Di Jin, Bo Yang, and Yuanfang Guo. Topology optimization based graph convolutional network. In *IJCAI*, pages 4054–4061, 2019.
- [42] X. Yang and G. E. Karniadakis. Reweighted ℓ_1 -minimization method for stochastic elliptic differential equations. *Journal of Computational Physics*, 248:87–108, 2013.
- [43] Yonggyun Yu, Taeil Hur, Jaeho Jung, and In Gwun Jang. Deep learning for determining a near-optimal topological design without any iteration. *Structural and Multidisciplinary Optimization*, 59(3):787–799, 2019.
- [44] Tomás Zegard and Glaucio H Paulino. Toward gpu accelerated topology optimization on unstructured meshes. *Structural and multidisciplinary optimization*, 48(3):473–485, 2013.
- [45] Yiquan Zhang, Airong Chen, Bo Peng, Xiaoyi Zhou, and Dalei Wang. A deep convolutional neural network for topology optimization with strong generalization ability. *arXiv preprint arXiv:1901.07761*, 2019.