

# DIFFUSING THE OPTIMAL TOPOLOGY: A GENERATIVE OPTIMIZATION APPROACH

Giorgio Giannone<sup>1,2,\*</sup>, Faez Ahmed<sup>1</sup>

<sup>1</sup>Massachusetts Institute of Technology, Cambridge, MA

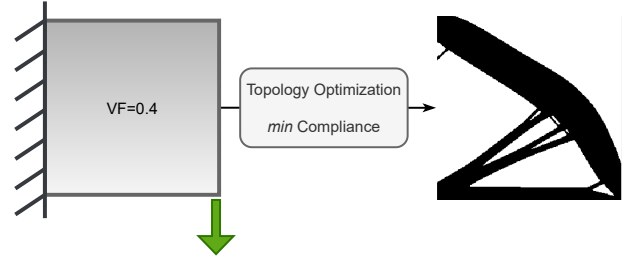
<sup>2</sup>Technical University of Denmark, Lyngby, DK

## ABSTRACT

*Topology Optimization seeks to find the best design that satisfies a set of constraints while maximizing system performance. Traditional iterative optimization methods like SIMP can be computationally expensive and get stuck in local minima, limiting their applicability to complex or large-scale problems. Learning-based approaches have been developed to accelerate the topology optimization process, but these methods can generate designs with floating material and low performance when challenged with out-of-distribution constraint configurations. Recently, deep generative models, such as Generative Adversarial Networks and Diffusion Models, conditioned on constraints and physics fields have shown promise, but they require extensive pre-processing and surrogate models for improving performance. To address these issues, we propose a Generative Optimization method that integrates classic optimization like SIMP as a refining mechanism for the topology generated by a deep generative model. We also remove the need for conditioning on physical fields using a computationally inexpensive approximation inspired by classic ODE solutions and reduce the number of steps needed to generate a feasible and performant topology. Our method allows us to efficiently generate good topologies and explicitly guide them to regions with high manufacturability and high performance, without the need for external auxiliary models or additional labeled data. We believe that our method can lead to significant advancements in the design and optimization of structures in engineering applications, and can be applied to a broader spectrum of performance-aware engineering design problems.*

## 1. INTRODUCTION

Topology optimization (TO [1, 2]) is an essential engineering tool that enables finding the optimal material distribution for meeting performance objectives while satisfying constraints. Constraints can take various forms, such as loads, boundary conditions, and volume fraction. Topology optimization can lead to significant improvements in the ef-



**Figure 1:** Minimum Compliance Topology Optimization problem. Given a set of constraints (loads, boundary conditions, volume fraction), find the optimal topology minimizing functional performance (structural compliance).

iciency, safety, reliability, and durability of structures across various fields, including aerospace, automotive, civil, and mechanical engineering. The demand for improvements in topology optimization techniques is ever-growing to enhance design engineering for a new age of engineering discovery. Topology optimization uses Finite Element Analysis (FEA), with gradient-based [1] and gradient-free methods [3] being two primary approaches. The Solid Isotropic Material with Penalization (SIMP) method [4, 5] is one of the most widely used gradient-based optimization algorithms. SIMP utilizes iterative optimization methods over continuous, density-based representation to determine the optimal material distribution. This class of methods iteratively adjusts the material density in a given domain to check whether the design meets the desired constraints.

Iterative optimization-based topology optimization (TO) methods have great potential benefits but face challenges in practical applications, particularly for large-scale problems due to their computational complexity. Iterative algorithms, like SIMP, are computationally expensive and prone to getting trapped in local optima, limiting the search for a global optimum. To address these challenges, there has been a recent surge of interest in learning-based methods for topology optimization, which use machine learning (ML) algorithms, such as deep generative models, to speed up the optimization

\*Corresponding author: ggiorgio@mit.edu, faez@mit.edu

process and generate more diverse structural topologies. By leveraging large datasets of existing designs, ML models can quickly identify patterns and generate new designs that meet specified constraints. Compared to traditional optimization methods, ML methods can handle high-dimensional data, explore a broader range of design solutions, and provide more diverse results that may not have been considered previously. Companies like SolidWorks and start-ups like nTopology are integrating deep generative topology optimization in their products, adding flexibility and creativity for the designers and speeding up the candidate generation phase.

**Deep Generative Models for TO.** Deep Generative models (DGMs [6]), have shown an impressive ability to model high-dimensional and complex data modalities, such as images or text, with high fidelity and diversity. This capability opens the door for significant improvements in creativity and productivity in various fields in the coming years. Recent advancements in large vision models [7]) and language models [8]) have greatly increased our capacity to process unstructured data, leading to the development of multimodal generation techniques [9]. These models hold great promise for enhancing engineering design [10, 11]. Building on these successes, DGMs have also been applied in the engineering field to problems with constraints, with a focus on improving the design process. However, these applications have primarily focused on metrics like reconstruction quality and have not fully addressed the fulfillment of engineering requirements.

DGMs for topology optimization pose challenges in terms of both performance and manufacturability, which TopologyGAN [12], addresses by conditioning on physics-based information. By introducing physical fields such as the Von Mises stress, strain energy density, and displacement fields, TopologyGAN improves the adherence of the generated samples to the underlying engineering problem. However, computing the physical fields requires solving a computationally expensive Finite Element Analysis (FEA) routine for each configuration of load and boundary condition considered, even though sampling new topologies with the base model (a conditional GAN) is fast. This approach has been successful but still optimizes for a pixel-wise, reconstruction-based loss that does not account for engineering requirements such as high performance and feasibility. Despite conditioning on physical fields, many designs generated by these models still suffer from floating materials that impede manufacturability, as well as limited diversity and generalizability out-of-distribution.

TopoDiff [13] proposes a conditional diffusion model, conditioning on fields similarly to [12], and an additional guidance mechanism to encourage the generative process to sample in regions with high manufacturability and high performance (low compliance in this specific case). This work solves some of the issues connected with performance and feasibility, at the cost of slow sampling, still relying on computationally expensive physical fields, and introducing surrogate models to account for performance and manufacturability. During inference, the model must compute the strain

and force fields for each constraint configuration using FEA and use this information to condition the model, enabling it to generate optimized topologies that meet the specified requirements. This conditioning step is crucial to ensure that the model performs well and generates output that is both feasible and optimized but, at the same time, expensive and not easy to scale to higher dimensionality, complex structures, and 3D domains.

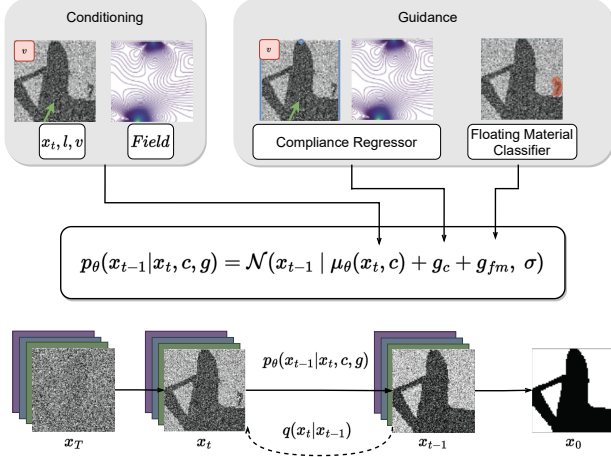
To summarize, GAN-based methods such as those proposed in [12, 14] can generate a large number of topologies efficiently but may produce un-manufacturable topologies that violate soft constraints such as volume fraction errors and lead to higher compliance structures. On the other hand, diffusion-based approaches like TopoDiff [13] generate samples that satisfy constraints more accurately but are computationally expensive due to iterative sampling, reliance on physical fields, and surrogate models with auxiliary labeled data.

To overcome the limitations of existing approaches, we propose a novel method that addresses the issues of slow sampling, reliance on physical fields, and the need for additional surrogate models. Our approach involves reducing the number of steps required for sampling, approximating physical fields using a computationally inexpensive kernel based on classic ODE solutions, and integrating optimization methods like SIMP for refining the generated topology. By explicitly guiding the generated topology to regions with high manufacturability and low compliance with only a few optimization steps, we can efficiently sample good topologies without the need for external auxiliary models, FEM solvers for pre-processing, or additional labeled data.

**Contribution.** Our contributions are the following:

- We decrease the computational cost of generative topology optimization while maintaining high performance. To achieve this, we explore more efficient sampling methods for TopoDiff reducing the number of steps required for generation by an order of magnitude while ensuring minimal loss in performance across all models.
- We explore alternative conditioning techniques that eliminate the need to compute force and energy strain fields, which can be a major bottleneck in the optimization process. As a result, we have reduced the inference time for generation by 53.93% compared to the baselines.
- We propose a generative optimization method that integrates a conditional diffusion model with traditional topology optimization algorithms. The fast conditional diffusion model predicts a preliminary topology, which is then refined using traditional topology optimization-based methods in just a few steps (5-10 iterations). This approach improves manufacturability and increases performance by 23.81% and 25.64%, respectively, for in- and out-of-distribution constraints.

## 2. RELATED WORK



**Figure 2: TopoDiff-GUIDED.** TopoDiff is a conditional diffusion model guided by a classifier and a regression score. The conditioning mechanism  $c$  involves loads  $l$ , volume fraction  $v$ , and stress and energy fields  $f$ . i.e.  $c = h(l, v, f)$ . The guidance mechanism involves a classifier for the presence or absence of floating material and a regressor to minimize compliance error between generated and optimized topologies.

**Topology Optimization.** Engineering design is the process of creating solutions to technical problems under engineering requirements [15]. The goal is to create designs that are highly performant given the required constraints. Topology Optimization (TO [1]) is a branch of engineering design and is a critical component of the design process in many industries, including aerospace, automotive, manufacturing, and software development. From the inception of the homogenization method for TO, a number of different approaches have been proposed, including density-based [4, 5], level-set [16], derivative-based [17], evolutionary [18], and others [19]. The density-based methods are widely used and use a nodal-based representation where the level-set leverages shapes derivative to obtain the optimal topology. Topology Optimization has evolved as a computationally intensive discipline, with the availability of efficient open-source code [20, 21]. See [21] for more on this topic and [2] for a comprehensive review of the Topology Optimization field.

**Deep Learning for Topology Optimization.** Following the success of Deep Learning (DL) in vision, a surging interest arose recently for transferring these methods to the engineering field. In particular, DL methods have been employed for direct-design [22], accelerating the optimization process [23], optimizing the shape [24], super-resolution [25, 26], and 3D topologies [27]. Among these methods, Deep Generative Models (DGMs) are especially appealing to improve design diversity in engineering design [28, 29]. Additionally, DGMs have been used for Topology Optimization problems conditioning on constraints (loads, boundary conditions, volume

fraction for the structural case), directly generating topologies [30, 31] training dataset of optimized topologies, leveraging superresolution methods to improve fidelity [32], using filtering and iterative design approaches [33] to improve quality and diversity. Methods for 3D topologies have also been proposed [14]. Recently, GAN-based approaches conditioning on constraints and physical information have had success in modeling the TO problem [12]. For a comprehensive review and critique of the field, see [34].

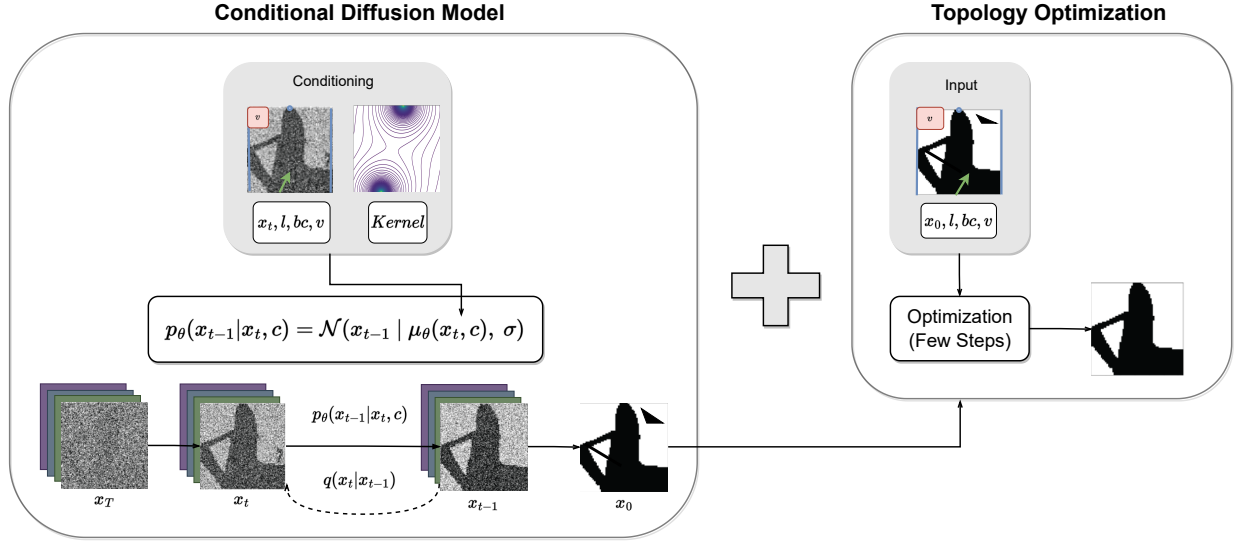
**Conditional Diffusion Models.** Improving sampling speed for diffusion models is an active research topic [35]. Recently, distillation has been used to greatly reduce sampling steps [36]. Methods to condition DDPM have been proposed, conditioning at inference time [37], learning a class-conditional score [38], explicitly conditioning on class information [39], set-based features [40], and physical properties [41]. Text-to-image diffusion models [42] have been recently proposed for a guided generation. Recently, TopoDiff [13] has shown that conditional diffusion models are effective for generating topologies that fulfill the constraints, and have high manufacturability and high performance. TopoDiff relies on physics information and surrogate models to guide the sampling of novel topologies with good performance.

## 3. BACKGROUND

Here we briefly introduce the Topology Optimization problem [1], diffusion models [43], a class of deep generative models, conditioning and guidance mechanisms for diffusion models, and recent deep generative models for topology optimization [12, 13].

**The Topology Optimization Problem.** Topology optimization (TO) is a powerful computational design approach used to determine the optimal configuration of a given structure, given a set of constraints. The goal of topology optimization is to identify the most efficient use of material while ensuring that the structure satisfies specific performance requirements. One common approach to topology optimization is the SIMP (Solid Isotropic Material with Penalization) method [4]. The SIMP method involves modeling the material properties using a density field, where the density represents the proportion of the material present in a given region. The optimization process involves adjusting the density field iteratively, subject to various constraints such as stress and boundary conditions. The objective of a generic minimum compliance problem for a mechanical system is to find the material density distribution  $\mathbf{x} \in \mathbb{R}^n$  that minimizes the structure’s deformation under the prescribed boundary and loads condition [2, 21]. Given a set of design variables  $\mathbf{x} = \{\mathbf{x}_i\}_{i=0}^n$ , where  $n$  is the domain dimensionality (in our case  $n = 64 \times 64$ ), the minimum compliance problems (Fig. 1) can be written as:

$$\begin{aligned} \min_{\mathbf{x}} \quad & c(\mathbf{x}) = F^T U(\mathbf{x}) \\ \text{s.t.} \quad & v(\mathbf{x}) = v^T \mathbf{x} < \bar{v} \\ & 0 \leq \mathbf{x} \leq 1 \end{aligned} \quad (1)$$



**Figure 3:** TopoDiff-FF + SIMP. The TopoDiff-FF pipeline is a conditional diffusion model where we condition on a cheap kernel relaxation instead of using expensive FEA to obtain the stress and energy field like in TopoDiff. After the generation step, we can improve the generated topology using few steps of SIMP (5/10 iterations) to remove floating material and explicitly optimize performance (minimize structural compliance).

	Load	BC	Kernel Load	Kernel BC	Force Field	Energy Field	VF	Performance
SIMP [1]	✓	✓	✗	✗	✗	✗	✓	✓
TopologyGAN [12]	✓	✓	✗	✗	✓	✓	✓	✗
TopoDiff [13]	✓	✓	✗	✗	✓	✓	✓	✗
TopoDiff-GUIDED [13]	✓	✓	✗	✗	✓	✓	✓	✓
TopoDiff-FF (ours)	✓	✓	✓	✓	✗	✗	✓	✗
TopoDiff-FF + SIMP (ours)	✓	✓	✓	✓	✗	✗	✓	✓

**Table 1:** Conditioning and guiding variables for different optimization methods and model configurations. All models condition directly or indirectly on loads, boundary conditions, and volume fraction. TopologyGAN, TopoDiff, and TopoDiff-GUIDED have additional conditioning on stress and energy fields, where TopoDiff-FF conditions are on the kernels. TopoDiff-GUIDED, TopoDiff-FF + SIMP, and SIMP itself are guided by a measure of performance.

where the goal is to find the design variables that minimize compliance  $c(\mathbf{x})$  given the constraints.  $F$  is the tensor of applied loads and  $U(\mathbf{x})$  is the node displacement, solution of the equilibrium equation  $K(\mathbf{x})U(\mathbf{x}) = F$  where  $K(\mathbf{x})$  is the stiffness matrix and is a function of the considered material.  $v(\mathbf{x})$  is the required volume fraction. The problem is a relaxation of the topology optimization task, where the design variables are continuous between 0 and 1. A simple thresholding mechanism is employed to assign the presence or absence of material in each node in the design domain. One significant advantage of topology optimization is its ability to create optimized structures that meet specific performance requirements. However, a major drawback of topology optimization is that it can be computationally intensive and may require significant computational resources. Additionally, some approaches to topology optimization may be limited in their ability to generate highly complex geometries and get stuck in local minima.

**Diffusion Models.** Let  $\mathbf{x}$  denote the observed data which is either continuous  $\mathbf{x} \in \mathbb{R}^D$  or discrete  $\mathbf{x} \in \{0, \dots, 255\}^D$ . Let  $\mathbf{z}_1, \dots, \mathbf{z}_T$  denote  $T$  latent variables in  $\mathbb{R}^D$ . We now introduce, the *forward or diffusion process*  $q$ , the *reverse or generative process*  $p_\theta$ , and the objective  $L$ . The forward or diffusion process  $q$  is defined as [44]:

$$q(\mathbf{z}_{1:T}|\mathbf{x}) = q(\mathbf{z}_1|\mathbf{x}) \prod_{t=2}^T q(\mathbf{z}_t|\mathbf{z}_{t-1}), \quad (2)$$

$$q(\mathbf{z}_t|\mathbf{z}_{t-1}) = \mathcal{N}(\mathbf{z}_t|\sqrt{1-\beta_t}\mathbf{z}_{t-1}, \beta_t I) \quad (3)$$

The beta schedule  $\beta_1, \beta_2, \dots, \beta_T$  is chosen such that the final latent image  $\mathbf{z}_T$  is nearly Gaussian noise. The generative or inverse process  $p_\theta$  is defined as:

$$p_\theta(\mathbf{x}, \mathbf{z}_{1:T}) = p_\theta(\mathbf{x}|\mathbf{z}_1)p(\mathbf{z}_T) \prod_{t=2}^T p_\theta(\mathbf{z}_{t-1}|\mathbf{z}_t), \quad (4)$$

$$p_\theta(\mathbf{z}_{t-1}|\mathbf{z}_t) = \mathcal{N}(\mathbf{z}_{t-1}|\mu_\theta(\mathbf{z}_t, t), \sigma_t^2 I), \quad (5)$$

where  $p(\mathbf{z}_T) = \mathcal{N}(\mathbf{z}_T|0, I)$ , and  $\sigma_t^2$  often is fixed (e.g. to  $\sigma_t^2 = \beta_t$ ). The neural network  $\mu_\theta(\mathbf{z}_t, t)$  is shared among all



time steps and is conditioned on  $t$ . The model is trained with a re-weighted version of the ELBO that relates to denoising score matching [45]. The negative ELBO  $L$  can be written as

$$\mathbb{E}_q \left[ -\log \frac{p_\theta(\mathbf{x}, \mathbf{z}_{1:T})}{q(\mathbf{z}_{1:T}|\mathbf{x})} \right] = L_0 + \sum_{t=2}^T L_{t-1} + L_T, \quad (6)$$

where  $L_0 = \mathbb{E}_{q(\mathbf{z}_1|\mathbf{x})} [-\log p(\mathbf{x}|\mathbf{z}_1)]$  is the likelihood term (parameterized by a discretized Gaussian distribution) and, if  $\beta_1, \dots, \beta_T$  are fixed,  $L_T = \mathbb{KL}[q(\mathbf{z}_T|\mathbf{x}), p(\mathbf{z}_T)]$  is a constant. The terms  $L_{t-1}$  for  $t = 2, \dots, T$  can be written as:

$$\begin{aligned} L_{t-1} &= \mathbb{E}_{q(\mathbf{z}_t|\mathbf{x})} \left[ \mathbb{KL}[q(\mathbf{z}_{t-1}|\mathbf{z}_t, \mathbf{x}) | p(\mathbf{z}_{t-1}|\mathbf{z}_t)] \right] \\ &= \mathbb{E}_{q(\mathbf{z}_t|\mathbf{x})} \left[ \frac{1}{2\sigma_t^2} \|\mu_\theta(\mathbf{z}_t, t) - \tilde{\mu}(\mathbf{z}_t, \mathbf{x})\|_2^2 \right] + C_t, \end{aligned} \quad (7)$$

where  $C_t = \frac{D}{2} \left[ \frac{\beta_t}{\sigma_t^2} - 1 + \log \frac{\sigma_t^2}{\beta_t} \right]$ . By further applying the reparameterization trick [46], the terms  $L_{1:T-1}$  can be rewritten as a prediction of the noise  $\epsilon$  added to  $\mathbf{x}$  in  $q(\mathbf{z}_t|\mathbf{x})$ . Parameterizing  $\mu_\theta$  using the noise prediction  $\epsilon_\theta$ , we can write

$$L_{t-1, \epsilon}(\mathbf{x}) = \mathbb{E}_{q(\epsilon)} \left[ w_t \|\epsilon_\theta(\mathbf{z}_t(\mathbf{x}, \epsilon)) - \epsilon\|_2^2 \right], \quad (8)$$

where  $w_t = \frac{\beta_t^2}{2\sigma_t^2 \alpha_t (1 - \alpha_t)}$ , which corresponds to the ELBO objective [47].

**Conditioning.** Similarly, introducing a conditioning variable  $\mathbf{c}$ , conditional modeling can be written as

$$p_\theta(\mathbf{x}, \mathbf{z}_{1:T}|\mathbf{c}) = p_\theta(\mathbf{x}|\mathbf{z}_1, \mathbf{c}) p_\theta(\mathbf{z}_T|\mathbf{c}) \prod_{t=2}^T p_\theta(\mathbf{z}_{t-1}|\mathbf{z}_t, \mathbf{c}), \quad (9)$$

and we can similarly write the per-layer loss with noise prediction as:

$$L_{t-1, \epsilon}^c(\mathbf{x}) = \mathbb{E}_{q(\epsilon)} \left[ w_t \|\epsilon_\theta(\mathbf{z}_t(\mathbf{x}, \epsilon), \mathbf{c}) - \epsilon\|_2^2 \right]. \quad (10)$$

We will now discuss different ways to condition the model.

**Guidance.** Guidance is used to improve sample fidelity vs mode coverage in conditional diffusion models post-training [48, 49]. Classifier guidance [48] is based on bayesian inversion and the fact that inverse problems can be tackled relatively easily in diffusion models. In particular, the goal of the guidance is to shift the model mean  $\mu_\theta$  closer to the target guidance  $\mathbf{c}$ . The general idea in classifier guidance is to leverage Bayes inversion

$$p(\mathbf{z}|\mathbf{g}) \propto p(\mathbf{g}|\mathbf{z})p(\mathbf{z}), \quad (11)$$

and notice that training the conditional model  $p(\mathbf{z}|\mathbf{g})$  is equivalent to training an unconditional model  $p(\mathbf{z})$  and a classifier  $p(\mathbf{g}|\mathbf{z})$ . Then the training signal will be a composition of unconditional and conditional scores for the model. Taking the gradients, we can see that classifier guidance is shifting the unconditional score for each class [48]. We can leverage

similar ideas for feasible vs unfeasible designs. With similar reasoning, regression guidance [13], aims to guide the sampling process towards configurations with low regression error as the target. In particular, the regression target for generating topologies is low compliance error between predicted (with a regression model) and real (obtained running a FEM solver) compliance. Guidance is useful to reduce high-compliance configurations generated by the model.

**TopoDiff.** Conditional diffusion models have been adapted for constrained engineering problems with performance requirements. TopoDiff [13] proposes to condition on loads, volume fraction, and fields similarly to [12] to learn a constrained generative model (Fig. 2). In particular, the generative model can be written as:

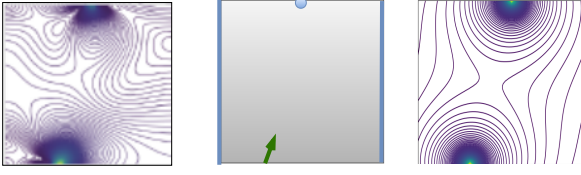
$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{c}, \mathbf{g}) = \mathcal{N}(\mathbf{x}_{t-1} | \mu_\theta(\mathbf{x}_t, \mathbf{c}) + \mathbf{g}_c + \mathbf{g}_{fm}, \sigma), \quad (12)$$

where  $\mathbf{c}$  is a conditioning term and is a function of the loads  $l$ , volume fraction  $v$ , and fields  $f$ , i.e  $\mathbf{c} = h(l, v, f)$ . The fields considered are the Von Mises stress  $\sigma_{vm} = (\sigma_{11}^2 - \sigma_{11}\sigma_{22} + \sigma_{22}^2 + 3\sigma_{12}^2)^{1/2}$  and the strain energy density field  $W = (\sigma_{11}\epsilon_{11} + \sigma_{22}\epsilon_{22} + 2\sigma_{12}\epsilon_{12})/2$ . Here  $\sigma_{ij}$  and  $\epsilon_{ij}$  are the stress and energy components over the domain.  $\mathbf{g}$  is a guidance term, containing information to guide the sampling process toward regions with low floating material (using a classifier and  $\mathbf{g}_{fm}$ ) and regions with low compliance error, where the generated topologies are close to optimized one (using a regression model and  $\mathbf{g}_c$ ). Where conditioning  $\mathbf{c}$  is always present and applied during training, the guidance mechanism  $\mathbf{g}$  is optional and applied only at inference time.

**Limitations.** TopoDiff is an effective method for generating topologies that satisfy constraints and have low compliance errors. However, the generative model is computationally expensive as hundreds of layers need to be sampled for each sample. Furthermore, the model requires preprocessing of configurations using a FEM solver, which is also computationally expensive and time-consuming. This approach relies heavily on fine-grained knowledge of the domain input, limiting its applicability to more challenging topology problems. Additionally, the model requires training of two surrogate models (a classification and regression model) which can be helpful for out-of-distribution configurations. However, to train the regression model, additional optimal and suboptimal topologies are needed assuming access to the desired performance metric on the train set. Similarly, the classifier requires additional labeled data to be gathered.

#### 4. METHOD

Our three main objectives are: improving inference efficiency, improving the sampling time for diffusion-based topology generation while still satisfying the design requirements with a minimum decrease in performance; minimizing reliance on force and strain fields as conditioning information, reducing the computation burden at inference time and the need for ad-hoc conditioning mechanisms for each



**Figure 4:** Field (left) vs Kernel (right) conditioning. Where computing the fields requires solving an expensive iterative FEA problem for each new configuration, computing the kernel relaxation is a single-step, computationally inexpensive approximation that does not rely on domain knowledge and scales to any resolution or domain structure.

Class	Metrics	Goal	Challenging?
Hard-constraint	Loads Disrespect	Feasibility	No
Hard-constraint	Floating Material	Manufacturability	Yes
Soft-constraint	Volume Fraction	Min Cost	No
Functional Performance	Compliance Error	Max Performance	Yes
Modeling Requirements	Sampling Time	Fast Inference	Yes

**Table 2:** Design and Modelling requirements for a constrained generative model for topology optimization. Our goal is to improve the requirements that are challenging to fulfill. In this work, we focus on improving Floating Material, reducing Compliance Error, and reducing Sampling Time.

problem and domain; merging together learning-based and optimization-based methods, refining the topology generated using a conditional diffusion model, improving the final solution in terms of manufacturability and performance. We aim to improve generative models with constraints and eliminate the need for costly FEM solutions as conditioning mechanisms. We refer to our approach as TopoDiff-FF, short for Topology-Diffusion-Fields-Free, and to our approach with optimization-based refinement as TopoDiff-FF + SIMP.

**Conditioning on Hard and Soft Constraints.** All models are subject to conditioning based on loads, boundary conditions, and volume fractions. In addition, TopoDiff and TopoDiff-GUIDED undergo conditioning based on force field and energy strain, while TopoDiff-FF is conditioned based on kernel relaxation, which is discussed in the following paragraph (see Fig. 4).

**Green’s Functions.** To improve the efficiency of diffusion-based topology generation and minimize reliance on force and strain fields, we aim to relax boundary conditions and loads by leveraging kernels as approximations for the way such constraints act on the domain. One possible choice of kernel structure is inspired by Green’s method [50], which defines integral functions that are solutions to the time-invariant Poisson’s Equation [51], a generalization of Laplace’s Equation for point sources excitations. Poisson’s Equation can be written as  $\nabla_x^2 f(x) = h$ , where  $h$  is a forcing term and  $f$  is a generic function defined over the domain  $\mathcal{X}$ . This equation governs many phenomena in nature, and a special case is a forcing part  $h = 0$ , which yields the Laplace’s Equation

formulation commonly employed in heat transfer problems. Green’s method is a mathematical construction to solve partial differential equations without prior knowledge of the domain. The solutions obtained with this method are known as Green’s functions [52]. While solutions obtained with this method can be complex in general, for a large class of physical problems involving constraints and forces that can be approximated with points, a simple functional form can be derived by leveraging the idea of source and sink. Consider a laminar domain (e.g., a beam or a plate) constrained in a feasible way. If a point source is applied to this domain (e.g., a downward force on the edge of a beam or on the center of a plate) in  $x_f$ , such force can be described using the Dirac delta function,  $\delta(x - x_f)$ . The delta function is highly discontinuous but has powerful integration properties. In particular  $\int f(x)\delta(x - x_f)dx = f(x_f)$  over the domain  $\mathcal{X}$ . The solution of the time-invariant Poisson’s Equation with point concentrated forces can be written as a Green’s function solution, where the solution depends only on the distance from the force application point. In particular:

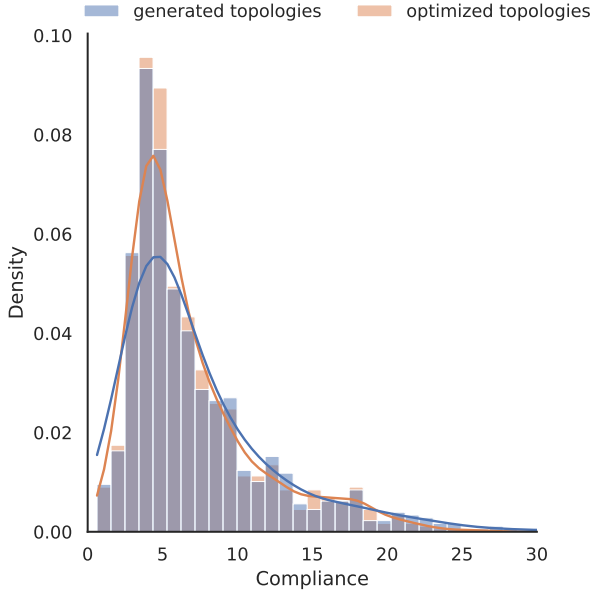
$$\mathcal{G}(x, x') = -\frac{1}{4\pi} \frac{1}{|x - x'|}, \quad (13)$$

where  $r = |x - x'| = \sqrt{|x_i - x'_i|^2 + |x_j - x'_j|^2}$ . We propose to approximate the forces and loads applied to our topologies using a kernel relaxation built using Green’s functions. While this formulation may not provide a correct solution for generic loads and boundary conditions, it allows us to provide computationally inexpensive conditioning information that respects the original physical and engineering constraints. By leveraging these ideas, we aim to increase the amount of information provided to condition the model, ultimately improving generative models with constraints.

**Kernel Relaxation.** We can use these kernels to construct a kernel relaxation method to condition generative models (Fig. 4). The idea is to use the kernels as approximations of the way boundary conditions and loads act on the domain. Specifically, we can use the kernels to represent the effects of the boundary conditions and loads as smooth functions across the domain. This approach avoids the need for computationally expensive and time-consuming finite element Analysis (FEA) to provide conditioning information. To apply the kernel relaxation method, we first determine the locations of the sources and sinks in the domain. Then, we compute the corresponding kernels for the loads and boundary conditions, respectively. Finally, we use the resulting kernels as conditioning information for generative models. In particular, we consider loads as sources and boundary conditions as sinks. For a load or source  $p$ , and  $r = |x - x^p| = \sqrt{|x_i - x_i^p|^2 + |x_j - x_j^p|^2}$  we have:

$$K_l(x, x^p; \alpha, \beta) \propto (1 - e^{-\alpha/r^\beta}) \bar{p}, \quad (14)$$

where  $\bar{p}$  is the module of a generic force in 2D. Notice how, for  $r \rightarrow 0$ ,  $K_l(x, x^p) \rightarrow p$ , and  $r \rightarrow \infty$ ,  $K_l(x, x^p) \rightarrow 0$ . For



**Figure 5:** Histogram empirical distribution generated and optimized compliances for task-2 (unknown constraints). We see that the generated topologies match well the distribution of compliance for unseen constraint configurations.

a boundary condition or sink:

$$K_{bc}(x, x^{bc}; \alpha, \beta) \propto e^{-\alpha/r^\beta}. \quad (15)$$

We notice how closer to the boundary the kernel is null, and farther from the boundary the kernel tends to 1. Note that the choice of  $\alpha$  and  $\beta$  parameters in the kernels affect the smoothness and range of the kernel functions. By adjusting these parameters, we can control the trade-off between accuracy and computational cost. Furthermore, these kernels are isotropic, meaning that they do not depend on the direction in which they are applied. Overall, the kernel relaxation method offers a computationally inexpensive way to condition generative models on boundary conditions and loads, making them more applicable in practical engineering and design contexts.

## 5. EXPERIMENTS

**Setup.** We train all the models for 200k steps on 30000 optimized topologies on a 64x64 domain. We set the hyperparameters, conditioning structure, and training routine as proposed in [13]. For all the models we condition on volume fraction and loads. For TopoDiff we condition on additional stress and energy fields. For TopoDiff-FF we condition on boundary conditions and kernels. TopoDiff-GUIDED leverages a compliance regressor and floating material classifier guidance (CLS+REG). We use a reduced number of sampling steps for all the experiments. See Table 2 for an overview of conditioning and guidance variables used for different models and methods considered in this work.

**Dataset.** We use the dataset of optimized topologies gathered using SIMP proposed in [13]. For each topology, we have

Cut-off	10	25	50	100
$F(x)$	0.9871	0.9963	0.9985	0.9990

**Table 3:** Cumulative Density Function for compliance cut-off on the train set. We can see that more than 99.90% of the training set has compliance lower than 100.

information about the loading condition, boundary condition, volume fraction, and optimal compliance. Additionally, for each configuration, a pre-processing step computes the force and strain energy fields (see Fig. 4).

**Evaluation.** We employ both engineering and generative metrics to evaluate the effectiveness of our model. Our metrics encompass physical, engineering, and modeling requirements to determine the model’s overall performance. We measure the Volume Fraction Error (VFE), which calculates the deviation between the generated topology volume and the prescribed volume. We also use Floating Material (FM) to assess the manufacturability of the model. For performance, we use the structural compliance metric [2], which measures the overall displacement under specified constraints. Minimizing compliance is one of the primary objectives of topology optimization. We calculate the average and median compliance (Avg C, Mdn C) for all tasks. We also use the Compliance Error (CE) to evaluate the engineering performance, which determines the deviation between the compliance calculated by the SIMP optimization method and the compliance generated by the diffusion model. Additionally, we analyze the inference time, which includes preprocessing and sampling time, to determine the model’s speed in creating fast concept designs. To assess the performance of the generative models, we design three test tasks: task-1, task-2, and task-3. We divide the test configurations into various groups based on their constraints and expected optimal performance. Table 4 presents these tasks. We notice that all the models have null load disrespect (LD), showing that fulfilling this hard constraint is easy for deep generative models.

- Task-1: the constraints in the task-1 test set are identical to those in the training set. We filter out generated configurations with high compliance when measuring performance on this task.
- Task-2: the constraints in the task-2 test set differ from those in the training set. We filter out generated configurations with high compliance when measuring performance on this task.
- Task-3: the constraints in the task-3 test set are the same as those in task-2. We consider all the generated configurations with low and high compliance when measuring performance on this task.

The purpose of these tasks is to evaluate the generalization capability of the machine learning models in- and out-of-distribution. By testing the models on different test sets with varying levels of difficulty, we can assess how well

the models can perform on new, unseen data. More importantly, we want to understand how important the role of the force field and energy strain is with unknown constraints.

task	constraints	performance	avg C	mdn C
task-1	in-distro	in-distro	7.73	5.54
task-2	out-distro	in-distro	6.84	5.29
task-3	out-distro	out-distro	6.85	5.31

**Table 4:** Tasks in order of challenge for the model. We expect the models to perform well on task-1, acceptably on task-2, and fail on task-3. We also report the average and median compliance over the different sets as a measure of optimality for the generated topologies.

### 5.1 Performance with In-Distribution Constraints

Table 5 presents an evaluation of different model variants on in-distribution constraints for Task 1, with various measures of compliance, compliance error, volume fraction error, and floating material. Lower is better for all the metrics. The models evaluated include TopoDiff, TopoDiff-GUIDED, TopoDiff-FF, and TopoDiff-FF + SIMP (10), with varying types of constraints and guidance. TopoDiff and TopoDiff-GUIDED condition of stress and energy fields, with TopoDiff-GUIDED leveraging surrogate models for guidance (classifier and regressor); TopoDiff-FF conditions on kernel relaxation. The table shows that TopoDiff and TopoDiff-GUIDED have lower values for most measures of compliance, compliance error, volume fraction error, and floating material compared to TopoDiff-FF. The TopoDiff-FF model has higher compliance error as expected but comparable VFE and FM, indicating that a lightweight kernel relaxation is a viable conditioning option. Additionally, TopoDiff-FF + SIMP (10) has the lowest values for compliance error, volume fraction error, and floating material. The table also reports the time taken for sampling, pre-processing, and inference for each model. TopoDiff-FF, even reducing the number of sampling steps for TopoDiff to 100, has significantly lower times for pre-processing and overall inference, because of the absence of FEA pre-processing. Overall, the results suggest that the TopoDiff and TopoDiff-GUIDED models perform better than TopoDiff-FF on in-distribution constraints, with TopoDiff-FF (and TopoDiff-FF+SIMP) being much faster at sampling and a viable data-driven alternative when FEA is impractical or computationally unfeasible. The results for TopoDiff-FF+SIMP suggest that using a combination of deep generative models and a simplified iterative method (SIMP) can lead to better results in terms of compliance and design requirements, even if the gap with TopoDiff is not so relevant when dealing with the in-distribution scenario.

### 5.2 Performance with Out-of-Distribution Constraints

Table 6 reports results on task-2 with out-of-distribution constraints. The considered models, conditioning, and guidance mechanisms are the same as explained in Subsec. 5.1. The results show that in out-of-distribution scenarios, TopoDiff-GUIDED model performs better than the

TopoDiff, achieving a lower CE and VFE. TopoDiff-FF performs well in terms of VFE and FM, but poorly in terms of compliance, with high CE. However, the TopoDiff-FF + SIMP (10) model significantly improves the performance on Task-2, achieving the lowest CE, VFE, and FM values among all models. As for task-1, the inference time for TopoDiff-FF and TopoDiff-FF+ SIMP is much better than for the full model. Overall, the table shows that the use of guidance and optimization techniques, such as SIMP, can significantly improve the performance of topology optimization models, especially on out-of-distribution constraints. See Fig 9 for a visual overview with confidence intervals over five runs.

### 5.3 Performance with Out-of-Distribution Constraints and Outliers

Table 7 presents results on task-3 with out-of-distribution and without filtering out outliers, i.e configurations with high-compliance (low performance). High-compliance configurations (>50/100) are out-of-distribution performance for the training set compliance distribution. In particular in Table 3 where we show that more than 99.8% of the configurations have compliance lower than 50. This means that when we obtain generated samples with such high compliance, the model has partially failed to fulfill the design constraints. The considered models, conditioning, and guidance mechanisms are the same as explained in Subsec. 5.1. The results in this table are similar to task-2, where TopoDiff-GUIDED performs better than TopoDiff in terms of CE and TopoDiff-FF achieves high compliance. However, using TopoDiff-FF+SIMP we not only improve all the metrics, but we completely get rid of such outlier configurations, being easy for the optimization process to fix the generated high-compliance configurations in a few iterations. Contrarily, for all other methods, there are always outliers in terms of topologies with high compliance that cannot easily be fixed, making a strong case for unifying deep generative models and optimization for topology optimization in challenging scenarios.

### 5.4 Compliance Analysis

Here we discuss and study how compliance influences the generated topology quality. Compliance magnitude seems to be the most important factor for the generation quality and manufacturability of proposed designs. In Fig. 5 we show qualitative information, providing empirical compliance distribution for generated vs optimized topologies. In Fig. 7 and Fig. 6 The growth of compliance error with compliance indicates that the model struggles to generate high-compliance topologies that fulfill the design requirements. This is further supported by the fact that very few configurations have high compliance, as indicated by the dataset. These findings highlight the difficulty of modeling compliant structures with high accuracy and precision and suggest that additional techniques or modifications to the model may be necessary to address this challenge. Moreover, the limited ability of the model to generate high-compliance topologies could have important implications for real-world applications, where compliant structures are often critical components of engineering



	Task	Steps	Constraints	Guidance	Avg C ↓	Mdn C ↓	% CE ↓	% VFE ↓	% FM ↓	Sampling ↓	Processing ↓	Inference ↓
TopoDiff	task-1	100	FIELD	COND	8.01	5.60	<u>5.46</u>	<u>1.47</u>	<u>5.79</u>	2.23	3.31	5.54
TopoDiff-GUIDED	task-1	100	FIELD	CLS+REG	8.15	5.63	5.93	1.49	5.82	2.46	3.31	5.87
TopoDiff-FF	task-1	100	KERNEL	COND	9.76	6.31	24.90	2.05	8.15	2.23	0.12	<b>2.35</b>
TopoDiff-FF + SIMP (10)	task-1	100	KERNEL	COND	8.05	5.77	<b>4.16</b>	<b>1.16</b>	<b>5.61</b>	2.23	0.32	<u>2.55</u>

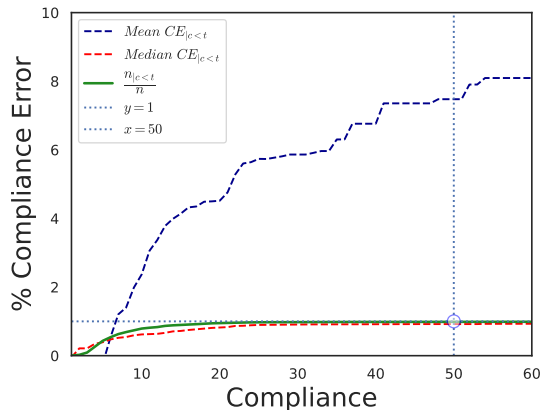
**Table 5:** Evaluation of different model variants on in-distribution constraints. We remove samples with compliance  $>100$ . C: compliance. CE: compliance error. VFE: volume fraction error. FM: floating material. We use 100 sampling steps for all models. SIMP (10) means that we run SIMP for 10 iterations. Results averaged over 5 runs. For a general overview with confidence intervals see the left side of Fig. 9.

	Task	Steps	Constraints	Guidance	Avg C ↓	Mdn C ↓	% CE ↓	% VFE ↓	% FM ↓	Sampling ↓	Processing ↓	Inference ↓
TopoDiff	task-2	100	FIELD	COND	7.80	5.49	12.02	1.49	<u>6.65</u>	2.23	3.31	5.54
TopoDiff-GUIDED	task-2	100	FIELD	CLS+REG	7.80	5.57	<u>10.55</u>	<u>1.47</u>	7.39	2.46	3.31	5.87
TopoDiff-FF	task-2	100	KERNEL	COND	11.65	6.94	58.36	1.97	7.86	2.23	0.12	<b>2.35</b>
TopoDiff-FF + SIMP (10)	task-2	100	KERNEL	COND	7.65	5.70	<b>7.84</b>	<b>1.29</b>	<b>6.53</b>	2.23	0.32	<u>2.55</u>

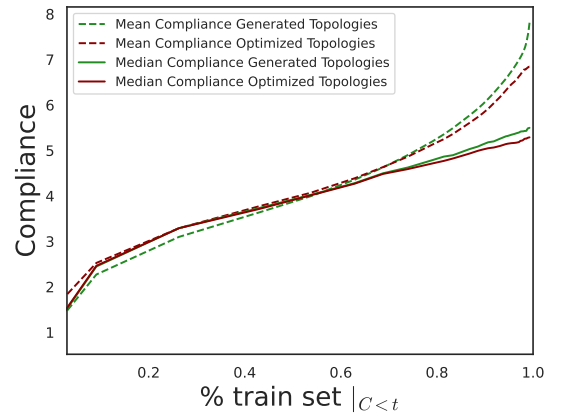
**Table 6:** Evaluation of different model variants on out-of-distribution constraints. We remove samples with compliance  $>100$ . C: compliance. CE: compliance error. VFE: volume fraction error. FM: floating material. We use 100 sampling steps for all models. SIMP (10) means that we run SIMP for 10 iterations. Results averaged over 5 runs. For a general overview with confidence intervals see the right side of Fig. 9.

designs. Therefore, improving the model’s ability to generate high-compliance topologies is an important research direction for advancing the state-of-the-art in this field.

**High-Compliance Configurations.** High-compliance configurations are often associated with mechanically unstable structures, meaning that they can collapse or deform easily under load. As a result, practical engineering designs often prioritize achieving low-compliance solutions that are stable, efficient, and lightweight. However, it is still important to study high-compliance configurations as they can provide valuable insights into the behavior and limitations of structural systems, as well as inspire new design concepts and approaches. We show some examples of high-compliance and high-compliance errors in Fig. 11. For low-compliance and low-compliance errors, see Fig. 12.



**Figure 6:** Mean and Median Compliance error for different compliance thresholds. We see that more than 99.8 % of the data has compliance lower than 50 and how the mean compliance error increases with compliance, where the median compliance error plateaus.



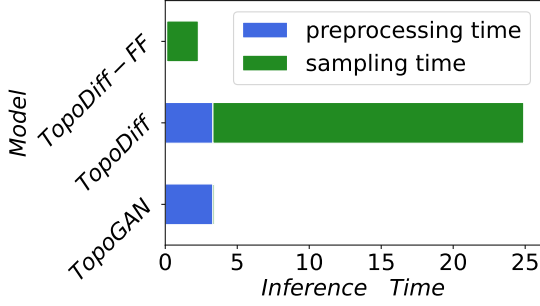
**Figure 7:** Mean and Median Compliance vs % train set for task-2 (unknown) constraints. The generated topologies using TopoDiff perform well for low compliance values. Most of the loss in terms of compliance happens for high compliance values.

## 5.5 Inference Time

Where in [13] has been shown that TopoDiff-based models largely outperform topologyGAN-based models [12] on design performances (CE, VFE, FM) for in- and out-distribution scenarios, inference time for GAN-based model is still much better than for Diffusion-based models. Here we present our approach to improving the efficiency of inference for TopoDiff-FF, measured in terms of both pre-processing and sampling time. Specifically, we focus on two key areas of optimization: I) Reducing the number of steps required for sampling to improve overall sampling time. II) Leveraging the kernel relaxation technique to reduce pre-processing time by avoiding the computation of force and energy fields. To demonstrate the effectiveness of our approach, we compare the pre-processing and sampling times of TopoGAN,

	Task	Steps	Constraints	Guidance	Avg C ↓	Mdn C ↓	% CE ↓	% VFE ↓	% FM ↓	Sampling ↓	Processing ↓	Inference ↓
TopoDiff	task-3	100	FIELD	COND	8.91	5.52	31.04	1.49	<u>6.61</u>	2.23	3.31	5.54
TopoDiff-GUIDED	task-3	100	FIELD	CLS+REG	8.23	5.58	<u>17.69</u>	<u>1.47</u>	7.36	2.46	3.31	5.87
TopoDiff-FF	task-3	100	KERNEL	COND	19.01	7.22	144.23	2.00	8.70	2.23	0.12	<b>2.35</b>
TopoDiff-FF + SIMP (10)	task-3	100	KERNEL	COND	7.65	5.70	<b>7.84</b>	<b>1.29</b>	<b>6.53</b>	2.23	0.32	<u>2.55</u>

**Table 7:** Evaluation of different model variants on out-of-distribution constraints. We do not filter out samples with high compliance. C: compliance. CE: compliance error. VFE: volume fraction error. FM: floating material. We use 100 sampling steps for all models. SIMP (10) means that we run SIMP for 10 iterations. Results averaged over 5 runs.



**Figure 8:** Sampling time. In this figure, we compare the total inference time (pre-processing and sampling time) for TopoDiff-FF, TopologyGAN, and TopoDiff-FF. We can see that we completely remove the pre-processing time component using the kernel relaxation. And reducing the number of sampling steps reduces, even more, the sampling time, making TopoDiff-FF relatively fast at the price of generating configurations with higher compliance than TopoDiff and TopoDiff-GUIDED.

TopoDiff, and TopoDiff-FF in Fig. 8. Our results show that TopoDiff-FF is able to generate high-quality samples that satisfy prescribed constraints and exhibit strong performance, all while requiring significantly less time for both pre-processing and sampling compared to TopoDiff and slightly less time than TopoGAN. This improved efficiency is a key goal of TopoDiff-FF, as it simplifies the conditioning process and enables faster inference.

## 5.6 Measuring the Relative Gap in Performance

To better summarize the results on in-distribution and out-of-distribution constraints, we propose to evaluate the models in terms of the relative design gap between the models, using a global metric that accounts for all the requirements presented in Table 2. We also consider relative ranking (Avg Rank) between the considered models. The gap and the rank are global metrics that account for compliance, compliance error, volume fraction error, floating material, processing, and inference time. For the average gap (Avg Gap), we use average compliance and compliance error. For the median gap (Mdn Gap), the corresponding median values. The design gap is a proxy to quantify the distance from an optimal design, where all the requirements are perfectly satisfied, we have compliance error null, and inference time is negligible. Table 2 presents a comparison of different conditioning configurations for generative designs in terms of their performance on task-1 and task-2. The table also shows whether the

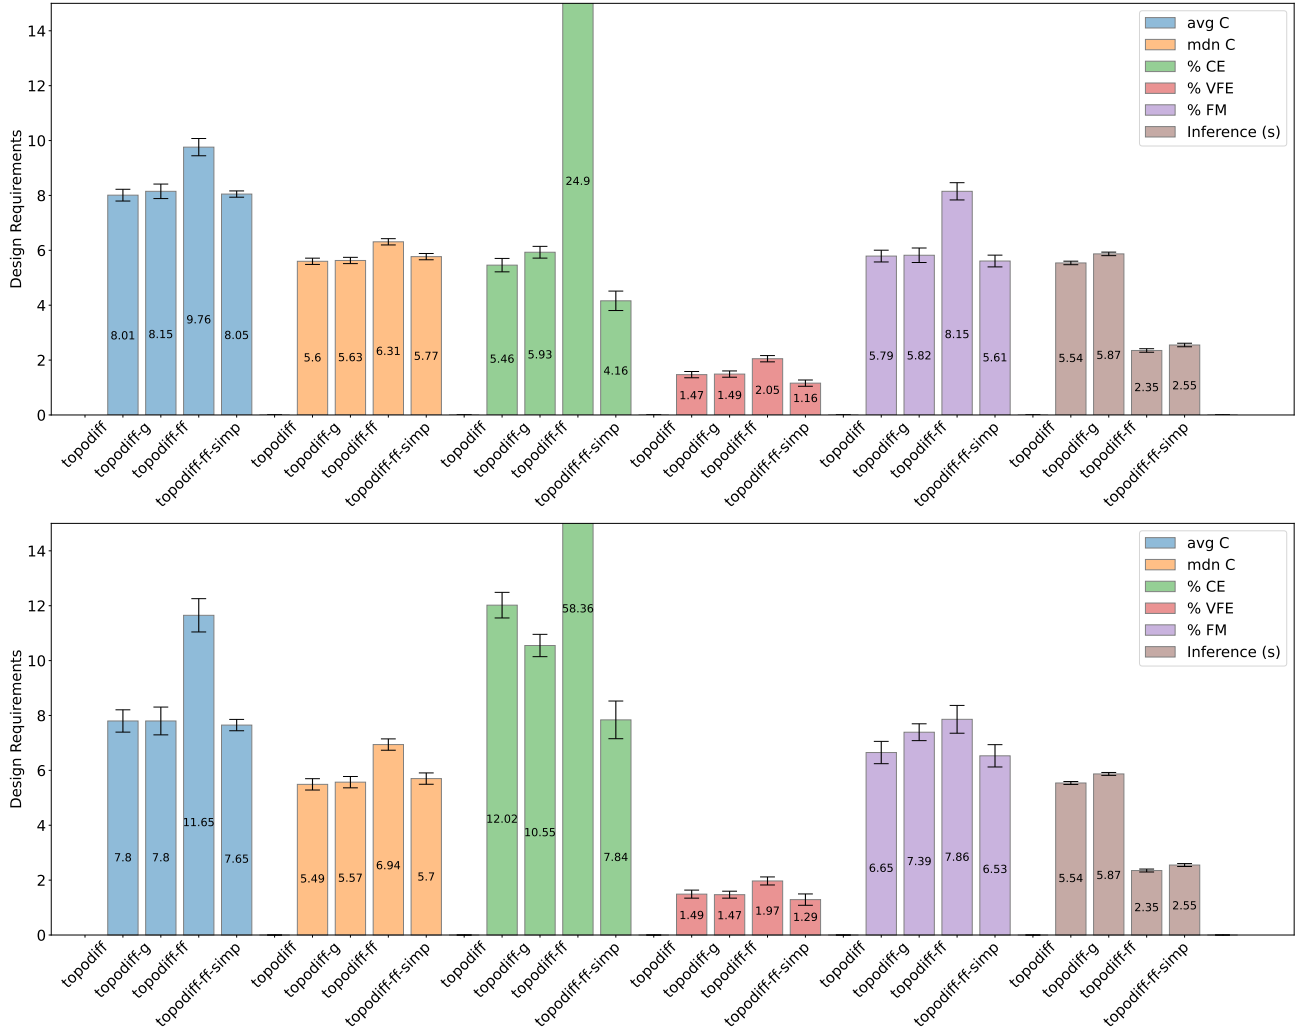
configuration includes kernels and field constraints. Overall, the models perform similarly in terms of the average and median gap on task-1, with TopoDiff-FF+SIMP (10) performing the best overall. On task-2, TopoDiff-FF loses in terms of the average gap because of the large CE on out-of-distribution. Interestingly, when considering average ranking, the gap is reduced because of the fast processing and inference time for TopoDiff-FF. As noted with the previous experiments, TopoDiff-FF with SIMP outperforms all other models in terms of both design quality and relative performance on both tasks, which suggests that the use of SIMP in conjunction with deep generative models is a promising approach for generating high-quality designs for topology optimization.

	Task	Kernels	Fields	Mdn Gap ↓	Avg Gap ↓	avg Rank ↓
TopoDiff	task-1	✗	✓	3.12	4.38	2.2
TopoDiff-GUIDED	task-1	✗	✓	3.16	4.52	3.2
TopoDiff-FF	task-1	✓	✗	4.37	7.86	2.8
TopoDiff-FF + SIMP	task-1	✓	✗	<b>2.69</b>	<b>3.57</b>	<b>1.6</b>
TopoDiff	task-2	✗	✓	3.35	5.58	2.2
TopoDiff-GUIDED	task-2	✗	✓	3.55	5.54	3.2
TopoDiff-FF	task-2	✓	✗	5.86	13.69	2.8
TopoDiff-FF + SIMP	task-2	✓	✗	<b>3.08</b>	<b>4.27</b>	<b>1.6</b>

**Table 8:** Global overview of different conditioning configurations showing how far the generative designs are from an ideal optimal design. On in-distribution configurations, our approach is extremely effective. On out-of-distribution configuration there is still a gap in performance. Adding few steps of SIMP (10 steps) as refinement is effective in improving the design quality.

## 5.7 Merging Deep Generative Models and Optimization

Table 9 shows the results of experiments conducted with three algorithms, SIMP, TopoDiff-FF, and TopoDiff-FF + SIMP, and tested on a particular data split. The table presents the values of two metrics, average C and Compliance Error, for each algorithm, with varying numbers of iterations. It is clear that increasing the number of iterations resulted in improved performance for all algorithms. Moreover, the combination of TopoDiff-FF and SIMP performed better than the other algorithms, achieving the lowest values of average C and CE in both experiments. These results suggest that using the combination of TopoDiff-FF and SIMP is a way to impose the performance constraints in the model without the need for surrogate models or guidance. Also, we are sure that some physical properties are respected, reducing the FM, VFE, and CE using a fast and relatively cheap refinement (5/10 iterations).



**Figure 9:** Overview design requirements for the considered models for task-1 (in-distribution constraints, top) and task-2 (out-of-distribution constraints, bottom). From left to right: average compliance in blue. Median compliance in orange. Compliance error in green. Volume fraction error in red. Floating material in purple. Inference time in brown. Lower is better for all the metrics.

	Task	Iter	Avg C ↓	% CE ↓
SIMP	task-2	+5	9.87	35.13
TopoDiff-FF	task-2	-	19.01	58.36
TopoDiff-FF + SIMP	task-2	+5	<b>8.67</b>	<b>20.34</b>
SIMP	task-2	+10	8.57	17.61
TopoDiff-FF	task-2	-	19.01	58.36
TopoDiff-FF + SIMP	task-2	+10	<b>7.65</b>	<b>7.84</b>

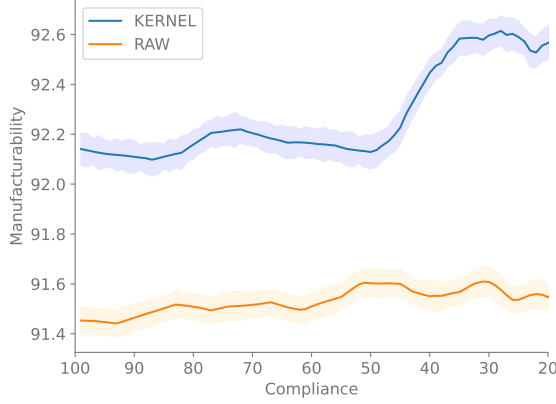
**Table 9:** Comparison of SIMP, TopoDiff-FF and TopoDiff-FF + SIMP. We see that by merging together the topology generated using a conditional diffusion model and SIMP for refinement we obtain the best results in terms of average compliance and compliance error with just 5 or 10 SIMP iterations.

### 5.8 Manufacturability

We define a topology as manufacturable if it satisfies all constraints, has no floating material, and compliance is below 100. In Figure 10, we depict how manufacturabil-

ity, measured as the percentage of configurations that meet these criteria, varies with performance. We evaluate out-of-distribution constraints and compare models trained on raw loads and boundary conditions with those trained on kernel relaxation. Our results show that manufacturability improves with performance, i.e., lower average compliance, of generated topologies. Notably, we observed a higher fraction of manufacturable topologies when conditioning on the kernel relaxation, supporting our hypothesis that the kernel is an effective approximation for loads and boundaries.

**Kernel Ablation.** The table presents the results of a kernel ablation study for the TopoDiff-FF model. We train the models using short runs. Four different inverse distance kernels with different exponents are evaluated, and their performance is measured in terms of compliance error (CE), volume fraction error (VFE), floating material (FM), load disrespect (LD), and the percentage of compliant samples with a compliance value smaller than 100 ( $conf|c < 100$ ). Lower values for



**Figure 10:** Manufacturability as measured in terms of absence of floating material and load disrespect, and compliance lower than 100 for out-of-distribution constraints.

CE, VFE, FM, and LD indicate better performance, while higher values for  $\text{conf}|c < 100$  indicate better compliance with the given constraints. The results show that the  $1/r^2$  kernel performs well overall, with the lowest CE and VFE values and the highest  $\text{conf}|c < 100$  percentage when learning the power using  $\beta$ . The  $1/r^\beta$  kernel also performs well, with the lowest FM value, while the  $1/r$  kernel performs the worst in all metrics.

	CE ↓	VFE ↓	FM ↓	LD ↓	$\text{conf} c < t \uparrow$
$K(1/r)$	189.76	15.54	8.60	1.4	48.80
$K(1/r^2)$	47.62	2.38	9.23	0.0	94.32
$K(1/r^4)$	55.01	2.35	8.15	0.0	95.90
$K(1/r^\beta)$	36.20	2.22	8.63	0.0	95.66

**Table 10:** Kernel Ablation for TopoDiff-FF. All metrics in %. We train models for a shorter time on a subset of the training data.

## 6. DISCUSSION

It is essential to acknowledge the limitations of purely data-driven approaches, such as TopoDiff-FF, which struggle to generalize to out-of-distribution scenarios since these models lack a way to generalize to new situations not encountered during training. In contrast, TopoDiff incorporates additional conditioning information obtained from numerical analysis and the underlying physics of the problem, which enables the model to generalize to regions of the domain not included in the training data. The kernel proposed for TopoDiff-FF is versatile and can be easily applied to any domain, resolution, and material, making it a valuable tool for engineering design. Contrary, FEM conditioning requires a thorough understanding of the problem, and it does not scale well with resolution and problem complexity. Therefore, TopoDiff-FF represents a powerful alternative for engineering design, particularly in situations where FEM conditioning is not feasible or effective. Furthermore, our results demonstrate that combining the topology generated by TopoDiff-FF with a few iterations

of SIMP can produce better results than either data-driven or physics-informed methods, indicating a way to increase expressivity for such models in engineering design.

**Limitations.** Deep Generative models for topology optimization are a promising direction to improve efficiency, scalability, and variety in engineering design but they still present limitations [34]. In general, three main objectives can be identified for data-driven approaches: direct design, speedup, and upsampling. The authors in [34] are critical of end-to-end approaches for direct design, stating that most of the methods proposed in the literature produce poor designs and are expensive and limited in the variety of problems and mesh resolutions they can handle. This is indeed a good critique because most methods focus on a specific domain and resolution. The authors argue that the iterative nature of topology optimization is not what makes it impractical but rather the computational load of the costly components within each iteration. They suggest that the focus should shift towards alleviating this computational load. The authors also suggest that the main contribution of the learning methods should be in speeding up the computations in the intermediate steps of the iterative optimization process or post-processing optimized results for manufacturability. Where most of these critiques are accurate, there is a new wave of generative models that are tackling such challenges and solving or alleviating most of these issues, from sampling efficiency to reliance on a specific domain; to optimization-based refinement and direct design with performance awareness [10].

## 7. CONCLUSION

In summary, our method of utilizing kernel approximation as conditioning for generative models and refining with optimization techniques has shown promise in enhancing precision, constraint satisfaction, manufacturability, and performance. To further improve our approach, future research could explore methods for maintaining diversity while preserving precision and constraint satisfaction. Furthermore, the scalability of our approach to larger and more complex problems should be investigated. In conclusion, we believe that our work offers new possibilities for advancing generative design in topology optimization and engineering design.

## REFERENCES

- [1] Bendsøe, Martin Philip and Kikuchi, Noboru. “Generating optimal topologies in structural design using a homogenization method.” *Computer methods in applied mechanics and engineering* Vol. 71 No. 2 (1988): pp. 197–224.
- [2] Sigmund, Ole and Maute, Kurt. “Topology optimization approaches: A comparative review.” *Structural and Multidisciplinary Optimization* Vol. 48 No. 6 (2013): pp. 1031–1055.
- [3] Ahmed, Faez, Bhattacharya, Bishakh and Deb, Kalyanmoy. “Constructive solid geometry based topology optimization using evolutionary algorithm.” *Proceedings of seventh international conference on bio-inspired computing: theories and applications (BIC-TA 2012)*: pp. 227–238. 2013. Springer.



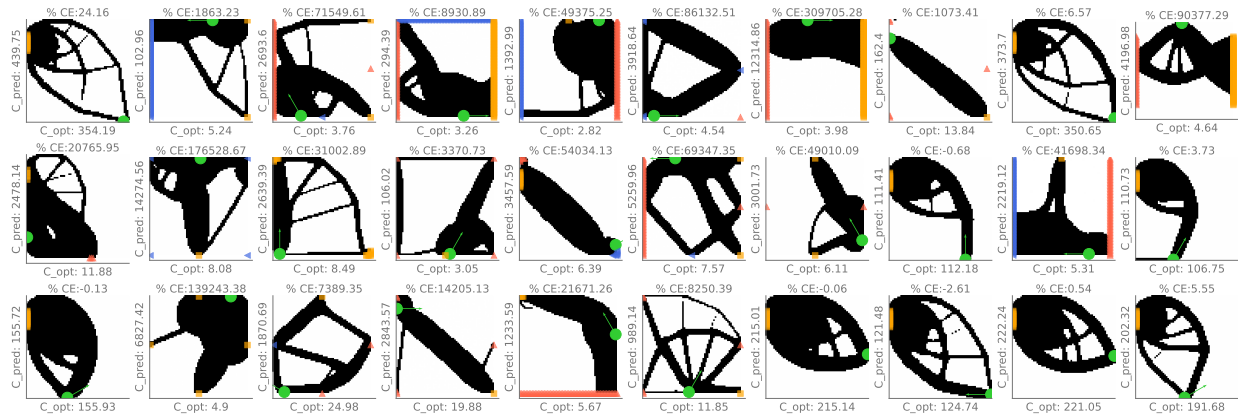


Figure 11: Examples of generated topologies with bad performance.

- [4] Bendsøe, Martin P. “Optimal shape design as a material distribution problem.” *Structural optimization* Vol. 1 (1989): pp. 193–202.
- [5] Rozvany, George IN, Zhou, Ming and Birker, Torben. “Generalized shape optimization without homogenization.” *Structural optimization* Vol. 4 (1992): pp. 250–252.
- [6] Bond-Taylor, Sam, Leach, Adam, Long, Yang and Willcocks, Chris G. “Deep generative modelling: A comparative review of vaes, gans, normalizing flows, energy-based and autoregressive models.” *IEEE transactions on pattern analysis and machine intelligence* (2021).
- [7] Rombach, Robin, Blattmann, Andreas, Lorenz, Dominik, Esser, Patrick and Ommer, Björn. “High-Resolution Image Synthesis with Latent Diffusion Models.” *arXiv preprint arXiv:2112.10752* (2021).
- [8] Brown, Tom B, Mann, Benjamin, Ryder, Nick, Subbiah, Melanie, Kaplan, Jared, Dhariwal, Prafulla, Neelakantan, Arvind, Shyam, Pranav, Sastry, Girish, Askell, Amanda et al. “Language models are few-shot learners.” *arXiv preprint arXiv:2005.14165* (2020).
- [9] Ramesh, Aditya, Pavlov, Mikhail, Goh, Gabriel, Gray, Scott, Voss, Chelsea, Radford, Alec, Chen, Mark and Sutskever, Ilya. “Zero-shot text-to-image generation.” *International Conference on Machine Learning*: pp. 8821–8831. 2021. PMLR.
- [10] Regenwetter, Lyle, Nobari, Amin Heyrani and Ahmed, Faez. “Deep generative models in engineering design: A review.” *Journal of Mechanical Design* Vol. 144 No. 7 (2022): p. 071704.
- [11] Song, Binyang, Zhou, Rui and Ahmed, Faez. “Multi-modal Machine Learning in Engineering Design: A Review and Future Directions.” *arXiv preprint arXiv:2302.10909* (2023).
- [12] Nie, Zhenguo, Lin, Tong, Jiang, Haoliang and Kara, Levent Burak. “Topologygan: Topology optimization using generative adversarial networks based on physical fields over the initial domain.” *Journal of Mechanical Design* Vol. 143 No. 3 (2021).
- [13] Maze, Francois and Ahmed, Faez. “TopoDiff: A Performance and Constraint-Guided Diffusion Model for Topology Optimization.” *arXiv preprint arXiv:2208.09591* (2022).
- [14] Behzadi, Mohammad Mahdi and Ilies, Horea T. “Real-Time Topology Optimization in 3D via Deep Transfer Learning.” *Computer-Aided Design* Vol. 135 (2021): p. 103014.
- [15] Shigley, Joseph Edward, Mitchell, Larry D and Saunders, H. “Mechanical engineering design.” (1985).
- [16] Allaire, Grégoire, Jouve, François and Toader, Anca-Maria. “A level-set method for shape optimization.” *Comptes Rendus Mathématique* Vol. 334 No. 12 (2002): pp. 1125–1130.
- [17] Sokolowski, Jan and Zochowski, Antoni. “On the topological derivative in shape optimization.” *SIAM journal on control and optimization* Vol. 37 No. 4 (1999): pp. 1251–1272.
- [18] Xie, Y Mike, Steven, Grant P, Xie, YM and Steven, GP. *Basic evolutionary structural optimization*. Springer (1997).
- [19] Bourdin, Blaise and Chambolle, Antonin. “Design-dependent loads in topology optimization.” *ESAIM: Control, Optimisation and Calculus of Variations* Vol. 9 (2003): pp. 19–48.
- [20] Hunter, William et al. “Topy-topology optimization with python.” (2017).
- [21] Liu, Kai and Tovar, Andrés. “An efficient 3D topology optimization code written in Matlab.” *Structural and Multidisciplinary Optimization* Vol. 50 (2014): pp. 1175–1196.
- [22] Abueidda, Diab W., Koric, Seid and Sobh, Nahil A. “Topology optimization of 2D structures with nonlinearities using deep learning.” *Computers & Structures* Vol. 237 (2020): p. 106283.
- [23] Banga, Saurabh, Gehani, Harsh, Bhilare, Sanket, Patel, Sagar and Kara, Levent. “3D Topology Optimization using Convolutional Neural Networks.” *Preprint* (2018).
- [24] Hertlein, Nathan, Buskohl, Philip R., Gillman, Andrew, Vemaganti, Kumar and Anand, Sam. “Generative adversarial network for early-stage design flexibility in topology optimization for additive manufacturing.” *Journal of Manufacturing Systems* Vol. 59 (2021): pp. 675–685.
- [25] Elingaard, Martin Ohrt, Aage, Niels, Bærentzen, Jakob Andreas and Sigmund, Ole. “De-homogenization using convolutional neural networks.” *Computer Methods in Applied Mechanics and Engineering* Vol. 388 (2022): p. 114197.
- [26] Napier, Nicholas, Sriraman, Sai-Akshar, Tran, Huy T. and James, Kai A. “An Artificial Neural Network Approach for Generating High-Resolution Designs From Low-Resolution Input in Topology Optimization.” *Journal of Mechanical Design* Vol. 142 No. 1 (2020).
- [27] Kench, Steve and Cooper, Samuel J. “Generating three-dimensional structures from a two-dimensional slice with generative adversarial network-based dimensionality expansion.” *Nature Machine Intelligence* Vol. 3 No. 4 (2021): pp. 299–305.
- [28] Jiang, Jiaqi, Chen, Mingkun and Fan, Jonathan A. “Deep neural networks for the evaluation and design of photonic devices.” *Nature Reviews Materials* Vol. 6 No. 8 (2021): pp. 679–700.

- [29] Rawat, Sharad and Shen, M. H. Herman. “A novel topology design approach using an integrated deep learning network architecture.” *Preprint* (2018).
- [30] Rawat, Sharad and Shen, MH Herman. “Application of adversarial networks for 3d structural topology optimization.” Technical report no. SAE Technical Paper. 2019.
- [31] Sharpe, Conner and Seepersad, Carolyn Conner. “Topology design with conditional generative adversarial networks.” *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, Vol. 59186: p. V02AT03A062. 2019. American Society of Mechanical Engineers.
- [32] Yu, Yonggyun, Hur, Taeil, Jung, Jaeho and Jang, In Gwun. “Deep learning for determining a near-optimal topological design without any iteration.” *Structural and Multidisciplinary Optimization* Vol. 59 No. 3 (2019): pp. 787–799.
- [33] Berthelot, David, Schumm, Thomas and Metz, Luke. “Began: Boundary equilibrium generative adversarial networks.” *arXiv preprint arXiv:1703.10717* (2017).
- [34] Woldseth, Rebekka V, Aage, Niels, Bærentzen, J Andreas and Sigmund, Ole. “On the use of artificial neural networks in topology optimisation.” *Structural and Multidisciplinary Optimization* Vol. 65 No. 10 (2022): p. 294.
- [35] Kong, Zhifeng and Ping, Wei. “On fast sampling of diffusion probabilistic models.” *arXiv preprint arXiv:2106.00132* (2021).
- [36] Meng, Chenlin, Gao, Ruiqi, Kingma, Diederik P, Ermon, Stefano, Ho, Jonathan and Salimans, Tim. “On distillation of guided diffusion models.” *arXiv preprint arXiv:2210.03142* (2022).
- [37] Choi, Jooyoung, Kim, Sungwon, Jeong, Yonghyun, Gwon, Youngjune and Yoon, Sungroh. “Ilvr: Conditioning method for denoising diffusion probabilistic models.” *arXiv preprint arXiv:2108.02938* (2021).
- [38] Song, Yang, Sohl-Dickstein, Jascha, Kingma, Diederik P, Kumar, Abhishek, Ermon, Stefano and Poole, Ben. “Score-based generative modeling through stochastic differential equations.” *arXiv preprint arXiv:2011.13456* (2020).
- [39] Nichol, Alexander Quinn and Dhariwal, Prafulla. “Improved denoising diffusion probabilistic models.” *International Conference on Machine Learning*: pp. 8162–8171. 2021. PMLR.
- [40] Giannone, Giorgio, Nielsen, Didrik and Winther, Ole. “Few-shot diffusion models.” *arXiv preprint arXiv:2205.15463* (2022).
- [41] Xie, Tian, Fu, Xiang, Ganea, Octavian-Eugen, Barzilay, Regina and Jaakkola, Tommi. “Crystal Diffusion Variational Autoencoder for Periodic Material Generation.” *arXiv preprint arXiv:2110.06197* (2021).
- [42] Ramesh, Aditya, Dhariwal, Prafulla, Nichol, Alex, Chu, Casey and Chen, Mark. “Hierarchical text-conditional image generation with clip latents.” *arXiv preprint arXiv:2204.06125* (2022).
- [43] Sohl-Dickstein, Jascha, Weiss, Eric, Maheswaranathan, Niru and Ganguli, Surya. “Deep Unsupervised Learning using Nonequilibrium Thermodynamics.” Bach, Francis and Blei, David (eds.). *Proceedings of the 32nd International Conference on Machine Learning*, Vol. 37: pp. 2256–2265. 2015. PMLR, Lille, France.
- [44] Ho, Jonathan, Jain, Ajay and Abbeel, Pieter. “Denoising Diffusion Probabilistic Models.” *Advances in Neural Information Processing Systems* 33. 2020.
- [45] Song, Yang and Ermon, Stefano. “Generative Modeling by Estimating Gradients of the Data Distribution.” *Advances in Neural Information Processing Systems* 32. 2019.
- [46] Kingma, Diederik P and Welling, Max. “Auto-encoding variational bayes.” *arXiv preprint arXiv:1312.6114* (2013).
- [47] Jordan, Michael I, Ghahramani, Zoubin, Jaakkola, Tommi S and Saul, Lawrence K. “An introduction to variational methods for graphical models.” *Machine learning* Vol. 37 No. 2 (1999): pp. 183–233.
- [48] Dhariwal, Prafulla and Nichol, Alexander. “Diffusion models beat gans on image synthesis.” *Advances in Neural Information Processing Systems* Vol. 34 (2021).
- [49] Ho, Jonathan and Salimans, Tim. “Classifier-Free Diffusion Guidance.” *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*. 2021.
- [50] Garabedian, PR. “Partial differential equations with more than two independent variables in the complex domain.” *Journal of Mathematics and Mechanics* (1960): pp. 241–271.
- [51] Hale, Jack K and Lunel, Sjoerd M Verduyn. *Introduction to functional differential equations*. Vol. 99. Springer Science & Business Media (2013).
- [52] Keldysh, Mstislav Vsevolodovich. “On the characteristic values and characteristic functions of certain classes of non-self-adjoint equations.” *Dokl. Akad. Nauk SSSR*, Vol. 77. 1: pp. 11–14. 1951.

## APPENDIX A. EXAMPLES OF GENERATED TOPOLOGIES



Figure 12: Examples of generated topologies with good performance.