

PHYSICS-INFORMED DIFFUSION MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Generative models such as denoising diffusion models are quickly advancing their ability to approximate highly complex data distributions. They are also increasingly leveraged in scientific machine learning, where samples from the implied data distribution are expected to adhere to specific governing equations. We present a framework that unifies generative modeling and **partial differential equation fulfillment by introducing a first-principle-based loss term that enforces generated samples to fulfill the underlying physical constraints. Our approach reduces the residual error by up to two orders of magnitudes compared to previous work in a fluid flow case study and outperforms task-specific frameworks in relevant metrics for structural topology optimization. We also present numerical evidence that our extended training objective acts as a natural regularization mechanism against overfitting.** Our framework is simple to implement and versatile in its applicability for imposing equality and inequality constraints as well as auxiliary optimization objectives.

1 INTRODUCTION

Denoising diffusion models (Sohl-Dickstein et al., 2015) are generative models that have gained popularity due to their success in learning intricate data distributions across various modalities, including images (Ho et al., 2020; Nichol et al., 2021; Rombach et al., 2021), videos (Ho et al., 2022b;a), graphs (Niu et al., 2020; Hoogetboom et al., 2022), text (Wu et al., 2023), and audio waveforms (Kong et al., 2020). Originally popularized within the image generation community, their outstanding representation abilities are also increasingly leveraged in the context of *scientific machine learning*. Diffusion models have been used, among others, to upscale low-fidelity data to reduce computational costs (Shu et al., 2023), design new molecules (Xu et al., 2022) or materials (Xie et al., 2021; Düreth et al., 2023) with desired properties, or generate metamaterial unit cells tailored to a given stress-strain response in complex mechanical settings (Bastek & Kochmann, 2023). For such inverse problems, a key strength is their ability to efficiently capture the full distribution of feasible solutions and designs rather than focusing on a single deterministic outcome.

A common thread of all those applications is the explicit knowledge of the underlying governing equations, which such implied distribution must obey. Often, the training data does not stem from experimental observations but is rather generated by numerical simulations, which ensure that those equations are fulfilled. Nevertheless, diffusion models are traditionally trained on a purely *data-driven* objective (Xie et al., 2021; Buehler, 2022; Bastek & Kochmann, 2023; Vlassis & Sun, 2023; Sardar et al., 2023; Li et al., 2023; Lienen et al., 2023) and hence do not strictly enforce the intrinsic constraints during model training. Consequently, samples generated from such models may be statistically aligned with the training data but may not meet the required precision in scientific applications, where adherence to the underlying physics is crucial for deployment (Jacobsen et al., 2023).

More recently, efforts have been made to ensure that samples generated from the learned distribution conform to the known constraints. Examples include imitating human motion (Yuan et al., 2022), ensuring manufacturability of proposed designs (Mazé & Ahmed, 2023), or, most importantly for physical systems, satisfying the underlying physical laws (Shu et al., 2023; Jacobsen et al., 2023), typically given by a set of partial differential equations (PDEs). Yet, a fundamental framework that rigorously addresses the incorporation of PDE constraints in such generative model settings, along with a robust mechanism to enforce these constraints akin to the well-established physics-informed neural networks (PINNs) (Raissi et al., 2019) have remained elusive. We close this gap by proposing

a new framework that unifies both settings and integrates PDE constraints meaningfully into the training process. Unlike previous approaches that primarily rely on some form of post-processing of generated samples and lack a derivation from first principles, we provide this theoretical foundation and directly embed constraints into the proven representation strength of diffusion models. This approach yields state-of-the-art results in terms of PDE fulfillment. By drawing synergies between the domains of PINNs (Raissi et al., 2019) and generative modeling, we introduce *physics-informed diffusion models* (PIDMs).

Contributions. We make the following key contributions: **(i)** We present a novel and rigorous approach that unifies denoising diffusion models with PINNs and informs the model of PDE constraints during training, and we demonstrate via rigorous numerical experiments that our framework significantly reduces the PDE residual compared to state-of-the-art methods. **(ii)** We provide evidence that the additional training objective does *not* necessarily compromise the data likelihood; instead, we observe that it acts as an effective regularization against overfitting. **(iii)** Our approach is simple to implement into the training protocol of existing diffusion model architectures, and inference is unaffected. **(iv)** While we here focus on PDEs as a sophisticated type of equality constraint, our framework is equally applicable to other equality and inequality constraints as well as auxiliary optimization objectives (potentially also provided via a differentiable surrogate model).

2 BACKGROUND

2.1 DENOISING DIFFUSION MODELS

Denoising diffusion models are state-of-the-art generative models (Ho et al., 2020; Dhariwal & Nichol, 2021; Yang et al., 2024) that learn to gradually convert a sample of a simple prior, typically a unit Gaussian, to a sample from a generally unknown data distribution $q(\mathbf{x}_0)$. The idea is to introduce a fixed *forward diffusion process* that incrementally adds Gaussian noise to a given data sample $\mathbf{x}_0 \sim q(\mathbf{x}_0)$, following variance schedule $\{\beta_t \in (0, 1)\}_{t=1}^T$ over T steps, defined as

$$q(\mathbf{x}_{1:T}|\mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t|\mathbf{x}_{t-1}), \quad q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I}). \quad (1)$$

To generate new samples, we consider the *reverse process*,

$$q(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^T q(\mathbf{x}_{t-1}|\mathbf{x}_t), \quad q(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}(\mathbf{x}_t, t), \boldsymbol{\Sigma}(\mathbf{x}_t, t)), \quad (2)$$

in which we approximate the unknown true inverse conditional distribution $q(\mathbf{x}_{t-1}|\mathbf{x}_t)$ with a neural network $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$ parameterized by θ . It aims to estimate the mean $\boldsymbol{\mu}_\theta$, while we fix the covariance to $\boldsymbol{\Sigma}(\mathbf{x}_t, t) = \frac{1 - \bar{\alpha}_t}{1 - \bar{\alpha}_t} \beta_t \mathbf{I} = \Sigma_t \mathbf{I}$ with $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$, $\alpha_t = 1 - \beta_t$. The network is trained by maximizing the variational lower bound of the log-likelihood, which can be simplified to several loss terms that mainly consist of KL-divergences between two Gaussians (and are hence computable in closed form) (Sohl-Dickstein et al., 2015). While the obvious choice is to estimate $\boldsymbol{\mu}_\theta$, alternative parameterizations are possible. We can obtain the mean $\boldsymbol{\mu}_t$ at timestep t via a combination of the reparameterization trick and Bayes’ theorem (Ho et al., 2020; Kingma & Welling, 2013):

$$\boldsymbol{\mu}_t(\mathbf{x}_t, \boldsymbol{\epsilon}_t) = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \boldsymbol{\epsilon}_t \right) = \frac{\sqrt{\bar{\alpha}_t} (1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} \mathbf{x}_t + \frac{\sqrt{\bar{\alpha}_{t-1}} \beta_t}{1 - \bar{\alpha}_t} \mathbf{x}_0, \quad (3)$$

where $\boldsymbol{\epsilon}_t$ represents Gaussian noise to diffuse \mathbf{x}_0 to \mathbf{x}_t . Since \mathbf{x}_t is known during training, predicting $\boldsymbol{\mu}_t$ fixes the Gaussian noise $\boldsymbol{\epsilon}_t$ or the clean signal \mathbf{x}_0 and vice versa, and we can equivalently train the model to predict these quantities. While Ho et al. (2020) simplified the training to minimize an unweighted noise mismatch, we here consider the loss

$$L(\theta) := \mathbb{E}_{t, \mathbf{x}_0, \boldsymbol{\epsilon}} \left[\lambda_t \|\mathbf{x}_0 - \hat{\mathbf{x}}_0(\mathbf{x}_t(\mathbf{x}_0, \boldsymbol{\epsilon}), t)\|^2 \right], \quad (4)$$

where $\hat{\mathbf{x}}_0$ is the model *estimate* of the clean signal (omitting the parametric dependence on θ for conciseness), and λ_t is set to Min-SNR-5 weighting (Hang et al., 2023). Note that equation 4 is equivalent to the mean squared error of the $\boldsymbol{\epsilon}_t$ or $\boldsymbol{\mu}_t$ mismatch up to a time-dependent weighting factor (Ho et al., 2020; Salimans & Ho, 2022).

2.2 ASSEMBLY OF GOVERNING EQUATIONS

Physical laws are typically formulated as a set of PDEs over a domain $\Omega \subset \mathbb{R}^d$, expressed as

$$\mathcal{F}[\mathbf{u}(\boldsymbol{\xi})] = \mathbf{0}, \quad \boldsymbol{\xi} = (\xi_1, \xi_2, \dots, \xi_d)^\top \in \Omega, \quad \mathbf{u} = (u_1(\boldsymbol{\xi}), u_2(\boldsymbol{\xi}), \dots, u_c(\boldsymbol{\xi}))^\top \in \mathbb{R}^c, \quad (5)$$

with boundary conditions

$$\mathcal{B}[\mathbf{u}(\boldsymbol{\xi})] = \mathbf{0}, \quad \boldsymbol{\xi} \in \partial\Omega, \quad (6)$$

where \mathcal{F} is a differential operator, \mathcal{B} is a boundary condition operator, $\partial\Omega$ is the boundary of domain Ω , and \mathbf{u} is the sought solution field that satisfies the set of PDEs for all $\boldsymbol{\xi} \in \Omega$ and boundary conditions on $\partial\Omega$. **We emphasize that we reserve $\boldsymbol{\xi}$ for spatial coordinates defined on a physical domain, while \mathbf{x} refers to data, which here typically contains the solution field on a discretized set of spatial coordinates, as elaborated below.**

We assume that samples generated from the model $\mathbf{x}_0 \sim p_\theta(\mathbf{x}_0)$ must satisfy equation 5 and equation 6. In this study, we mainly consider image-based architectures common in diffusion models. More formally, we consider a discretized pixel grid $\Omega^h \subset \mathbb{Z} \times \mathbb{Z}$ (which may serve as an approximation of the continuous domain \mathbb{R}^2), where $\partial\Omega^h$ consists of the boundary pixels. For $n \times n$ pixels, the model’s output is thus defined over $\mathbf{x}_0 \in \mathbb{R}^{c \times n \times n}$. For instance, in the context of Darcy flow problems in porous media (Jacobsen et al., 2023) \mathbf{x}_0 describes the pressure field, or in the context of solid mechanics the displacement field (Gao et al., 2022). While PINNs (Raissi et al., 2019) establish an explicit map $\mathcal{N} : \boldsymbol{\xi} \mapsto \mathbf{u}$ between the input and solution field, which enables the computation of required derivatives for \mathcal{F} (and potentially \mathcal{B}) via automatic differentiation, we here approximate those derivatives via finite differences or comparable methods that directly operate on \mathbf{x}_0 . The residual $\mathcal{R}(\mathbf{x}_0)$ of a solution field \mathbf{x}_0 is then established by concatenating $\mathcal{F}[\mathbf{x}_0]$ and $\mathcal{B}[\mathbf{x}_0]$, as presented in detail in A.9. For evaluation, we introduce the mean absolute error $\mathcal{R}_{\text{MAE}}(\mathbf{x}_0)$ as defined in equation 27 in Appendix A.6.

3 PHYSICS-INFORMED DIFFUSION MODELS

We explore a scenario in which our generative diffusion model must learn a distribution whose samples are to comply with a set of governing equations, i.e., $\mathcal{R}(\mathbf{x}_0) = \mathbf{0}$. In the usual data-driven setting, this is only indirectly ensured through the training data, typically collected from a forward simulator that produces data points which inherently follow the governing equations. These fully describe the physical system and provide us with, in principle, an *infinite* source of information (Rixner & Koutsourelakis, 2021), whereas solved instances $\{\mathbf{x}_0^1, \dots, \mathbf{x}_0^n\}$, posing as training data, represent only a finite set of evaluations in terms of the sought solution fields.¹ Thus, our strategy is to first state our optimization objective based on the more fundamental governing equations and only subsequently incorporate training data.

3.1 CONSIDERATION OF PDE CONSTRAINTS

We maintain the probabilistic perspective of generative models and introduce the residuals as *virtual observables* (Rixner & Koutsourelakis, 2021) $\hat{\mathbf{r}} = \mathbf{0}$, which we consider to be sampled from the distribution

$$q_{\mathcal{R}}(\hat{\mathbf{r}}|\mathbf{x}_0) = \mathcal{N}(\hat{\mathbf{r}}; \mathcal{R}(\mathbf{x}_0), \sigma^2 \mathbf{I}). \quad (7)$$

In the limit $\sigma^2 \rightarrow 0$, this recovers the deterministic setting of strictly enforcing the residual equations, as all probability mass is concentrated in $\mathcal{R}(\mathbf{x}_0)$. This can be used to compute the virtual likelihood $p_\theta(\hat{\mathbf{r}})$, which we expand in terms of drawn samples \mathbf{x}_0 as

$$p_\theta(\hat{\mathbf{r}}) = \int p_\theta(\hat{\mathbf{r}}, \mathbf{x}_0) d\mathbf{x}_0 = \int q_{\mathcal{R}}(\hat{\mathbf{r}}|\mathbf{x}_0) p_\theta(\mathbf{x}_0) d\mathbf{x}_0 = \mathbb{E}_{\mathbf{x}_0 \sim p_\theta(\mathbf{x}_0)} q_{\mathcal{R}}(\hat{\mathbf{r}}|\mathbf{x}_0). \quad (8)$$

This factorization of $p_\theta(\hat{\mathbf{r}}, \mathbf{x}_0)$ is reasonable as the residual follows \mathbf{x}_0 via equation 7. The goal is to maximize the usual log-likelihood over the virtual observable $\hat{\mathbf{r}}$, i.e.,

$$\arg \max_{\hat{\mathbf{r}}} \mathbb{E}_{\hat{\mathbf{r}}} [\log p_\theta(\hat{\mathbf{r}})] = \arg \max_{\hat{\mathbf{r}}} \mathbb{E}_{\mathbf{x}_0 \sim p_\theta(\mathbf{x}_0)} [\log q_{\mathcal{R}}(\hat{\mathbf{r}} = \mathbf{0}|\mathbf{x}_0)]. \quad (9)$$

¹Even if we assume access to infinite training samples, the residual information may still be beneficial, since it typically imposes constraints on higher-order derivatives (comparable to Sobolev training (Czarnecki et al., 2017)).

Thus, any samples from $p_\theta(\mathbf{x}_0)$ are evaluated on their log-likelihood via equation 7, which can also be understood as a probabilistic reinterpretation of the loss used to train PINNs (Raissi et al., 2019). Alternatively, equation 9 can also be understood as sampling from an unnormalized target distribution assembled by evaluating $\log q_{\mathcal{R}}$ over \mathbf{x}_0 , as is often considered in data-free settings (Zhang & Chen, 2022b; Vargas et al., 2023; Berner et al., 2024). Our setting, however, differs as we consider the more common scenario of applying diffusion models to learn a data distribution, as elaborated as follows.

3.2 CONSIDERATION OF OBSERVED DATA

We emphasize our focus on a *generative* model class in which training data is typically available. This is crucial for two main reasons: first, identifying solutions \mathbf{x}_0 that satisfy the governing equations is non-trivial, and we may understand the obtained training data as guidance for the model towards feasible solutions, which may accelerate the training. Alternatively, the constraints may be far from a “well-posed” problem and admit a large set of solutions, while we are only interested in a small subset of those reflected by the data. Second, equation 9 contains any distribution that produces samples with zero residuals, and the model may simply collapse to a single solution instance (such as in the classical PINN setup), which is not the use case for a generative scenario. Rather, we aim to leverage the proven capabilities of generative models to learn complex distributions, such as optimal mechanical designs or fluid flows based on some conditioning, while ensuring adherence to the physical laws. Even if no data is available, we may introduce an uninformative prior such as a unit Gaussian to promote the exploration of different solutions, which we briefly investigate in Appendix A.6.1.

We thus extend the objective equation 9 by including the usual data likelihood term, which is taken as the standard optimization objective in data-driven diffusion models as

$$\arg \max_{\theta} \mathbb{E}_{\mathbf{x}_0 \sim q(\mathbf{x}_0)} [\log p_\theta(\mathbf{x}_0)] + \mathbb{E}_{\mathbf{x}_0 \sim p_\theta(\mathbf{x}_0)} [\log q_{\mathcal{R}}(\hat{\mathbf{r}} = \mathbf{0} | \mathbf{x}_0)]. \quad (10)$$

We show in Appendix A.1 that, if diffusion models are interpreted as score-based models, a straightforward inclusion of the virtual (residual) likelihood in the loss function recovers a consistent training objective in the sense that the optimal score model also maximizes the introduced virtual (residual) likelihood and recovers samples from $q(\mathbf{x}_0)$.

3.3 SIMPLIFICATION OF THE TRAINING OBJECTIVE

The joint loss equation 10 requires sampling not just from $q(\mathbf{x}_0)$ but also $p_\theta(\mathbf{x}_0)$, which is costly for diffusion models due to their iterative nature. We explore two ideas to mitigate this increased complexity by considering (i) a straightforward evaluation of the residual based on the readily available $\hat{\mathbf{x}}_0$ and (ii) an accelerated sampling strategy via denoising diffusion implicit models (DDIMs) (Song et al., 2020).

Mean estimation. As seen in equation 4, the diffusion model objective can be understood as minimizing a (time-weighted) mean-squared distance between the predicted and true sample, given a noisy input. Hence, $\hat{\mathbf{x}}_0$ is an estimate of $\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t]$ (Song et al., 2021). Evaluating the residual on this estimate is therefore in general not fully consistent, as $\mathcal{R}(\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t]) \neq \mathbb{E}[\mathcal{R}(\mathbf{x}_0 | \mathbf{x}_t)]$ (also referred to as the *Jensen gap* (Gao et al., 2017)). Only at the last sampling step will this estimate align with a generated sample \mathbf{x}_0 (i.e., at $t = 1$, assuming $\beta_1 = 0$) and otherwise introduce a conflicting objective between the data and residual loss with increasing t . To mitigate this, we propose to increase the variance of the introduced residual likelihood with increasing t , while enforcing tighter adherence to the constraint as $t \rightarrow 0$.

Sample estimation. Alternatively, we may aim to evaluate the residual of an actual sample $\mathbf{x}_0 \sim p_\theta(\mathbf{x}_0)$, motivated by the implied consistency shown in Appendix A.1 (though in idealized settings). This comes at the expense of increased computational complexity due to the many forward passes required in the denoising process. We hence consider (deterministic) DDIM (Song et al., 2020) to accelerate sampling while maintaining high sample quality, with an inherent trade-off between sample quality and the number of DDIM timesteps. We consider a simple two-step sampling of \mathbf{x}_0 from any t and again increase the variance as $t \rightarrow T$, where timesteps are coarser, which leads to reduced sample quality (see Appendix A.4 for further details). While we focus on diffusion models, we note

that single-step generation models such as consistency models (Song et al., 2023) are a promising alternative, offering direct access to a sample.

Since the remaining derivations hold true for both the mean and sample estimation, we no longer distinguish between these estimates and denote both of them with \mathbf{x}_0^* , based on which the residual can efficiently be evaluated. As \mathbf{x}_0^* is obtainable at any timestep t , we adapt equation 10 to

$$\arg \max_{\theta} \mathbb{E}_{\mathbf{x}_0 \sim q(\mathbf{x}_0)} [\log p_{\theta}(\mathbf{x}_0)] + \mathbb{E}_{\mathbf{x}_{1:T} \sim p_{\theta}(\mathbf{x}_{1:T})} [\log q_{\mathcal{R}}(\hat{\mathbf{r}} = \mathbf{0} | \mathbf{x}_0^*(\mathbf{x}_{1:T}))]. \quad (11)$$

Estimates at noisier inputs may be less accurate for aforementioned reasons, thus we penalize constraint violation less as $t \rightarrow T$ by modeling $q_{\mathcal{R}}(\hat{\mathbf{r}} | \mathbf{x}_0^*(\mathbf{x}_{1:T}))$ for simplicity based on a scaled version of the variance scheduler of the reverse process as

$$q_{\mathcal{R}}(\hat{\mathbf{r}} | \mathbf{x}_0^*(\mathbf{x}_{1:T})) = \prod_{t=1}^T q_{\mathcal{R}}(\hat{\mathbf{r}} | \mathbf{x}_0^*(\mathbf{x}_t, t)), \quad \text{where } q_{\mathcal{R}}(\hat{\mathbf{r}} | \mathbf{x}_0^*(\mathbf{x}_t, t)) = \mathcal{N}(\hat{\mathbf{r}}; \mathcal{R}(\mathbf{x}_0^*), \Sigma_t/c \mathbf{I}), \quad (12)$$

where Σ_t is, as introduced before, the fixed variance of the denoising process. The scale factor $c > 0$ that is the sole hyperparameter of our setup effectively dictates the penalty of deviating from $\mathcal{R}(\mathbf{x}_0^*) = \mathbf{0}$. Figure 1 shows a graphical illustration of this process.

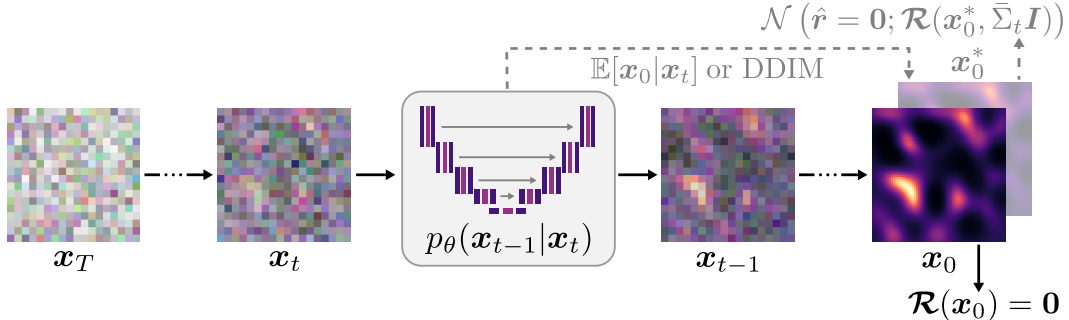


Figure 1: An approximation \mathbf{x}_0^* of the clean signal for residual evaluation can be obtained at any denoising timestep t by directly considering the estimated expectation $\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t]$ or by actual (accelerated DDIM (Song et al., 2020)) sampling. We tighten the variance of the virtual likelihood as $t \rightarrow 0$.

As equation 11 still requires sampling over $\mathbf{x}_{1:T} \sim p_{\theta}(\mathbf{x}_{1:T})$, we simplify this by instead sampling from the available $q(\mathbf{x}_{1:T})$, effectively ignoring the likelihood ratio (see Appendix A.3 for details). As the original (Sohl-Dickstein et al., 2015) and physics-driven objectives are then under the same expectation, the classical data loss equation 4 can be straightforwardly extended to include a physics-informed loss, resulting in the loss of our proposed PIDM:

$$L_{\text{PIDM}}(\theta) = \mathbb{E}_{t \sim [1, T], \mathbf{x}_{0:T} \sim q(\mathbf{x}_{0:T})} \left[\lambda_t \|\mathbf{x}_0 - \hat{\mathbf{x}}_0(\mathbf{x}_t, t)\|^2 + \frac{1}{2\Sigma_t} \|\mathcal{R}(\mathbf{x}_0^*(\mathbf{x}_t, t))\|^2 \right], \quad (13)$$

where $\bar{\Sigma}_t = \Sigma_t/c$ is the rescaled variance. Since $\Sigma_0 = 0$ for a deterministic last step, we set $\Sigma_0 \leftarrow \Sigma_1$. Algorithm 1 displays the updated training objective, requiring only minor modifications to the standard training setup (Ho et al., 2020).

We point out that the model can also be trained to match the mean or noise as in Ho et al. (2020), from which we can equally assemble $\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t]$ if considering mean estimation via equation 3 (see Appendix A.2 for estimation with score models) but we empirically found that this leads to larger residual errors. Also, we emphasize that PDEs can be understood as a specific instance of equality constraints and we could generally consider any differentiable forward surrogate model $\mathcal{R}_{\theta}(\mathbf{x}_0)$ that estimates some property or classification of the predicted samples to be matched, i.e., $\mathcal{R}_{\theta}(\mathbf{x}_0) = \hat{\mathbf{r}}_{\text{target}}$. For a discussion on how to incorporate inequality constraints and auxiliary optimization objectives see Appendix A.5. Lastly, we note that the optimal scaling c generally depends on the considered scenario, and it is here estimated by a simple parameter sweep. As usual in multi-objective optimization, trade-offs between the different loss contributions are expected, and c must be selected so that the model is meaningfully informed by the residual loss but does not ignore the data likelihood.

Algorithm 1 Physics-informed diffusion model training

-
- 1: Set $\bar{\Sigma}_t = \Sigma_t/c$
 - 2: **repeat**
 - 3: $\mathbf{x}_0 \sim q(\mathbf{x}_0)$
 - 4: $t \sim \text{Uniform}\{1, \dots, T\}$
 - 5: $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 - 6: $\mathbf{x}_t = \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\boldsymbol{\epsilon}$
 - 7: Estimate \mathbf{x}_0^* via $\hat{\mathbf{x}}_0 = \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t]$ or DDIM sampling (Song et al., 2020)
 - 8: Take gradient descent step on $\nabla_{\theta} \left[\lambda_t \|\mathbf{x}_0 - \hat{\mathbf{x}}_0(\mathbf{x}_t, t)\|^2 + \frac{1}{2\bar{\Sigma}_t} \|\mathcal{R}(\mathbf{x}_0^*(\mathbf{x}_t, t))\|^2 \right]$
 - 9: **until** converged
-

4 EXPERIMENTS

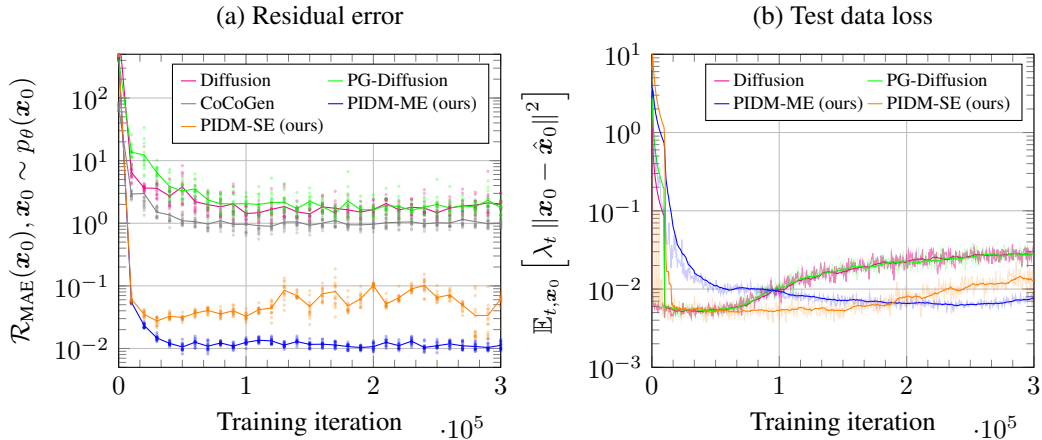
We here present two benchmarks for distributions from which samples are implied to adhere to specific governing PDEs and constraints, as recently studied in contemporary work (Jacobsen et al., 2023; Giannone et al., 2023). To obtain intuition for the effects of the proposed loss, see the toy problem presented in Appendix A.6.1 as an instructive example.

4.1 DARCY FLOW

Setup. We first study the 2D Darcy flow equations, which describe the steady-state solution for fluid flow through a porous medium, here on a square domain $\Omega = [0, 1]^2$. Generally, we follow the setup of previous studies (Jacobsen et al., 2023; Zhu & Zabaras, 2018) by sampling permeability fields $K(\boldsymbol{\xi})$ from a Gaussian random field, which are solved for their (unique) pressure distribution $p(\boldsymbol{\xi})$. Similarly, we create a training and a validation dataset of 10,000 and 1,000 datapoints, respectively, by solving the governing equations (see equation 28) for a sampled permeability field on a 64×64 grid. Second-order central finite differences are used to assemble and solve a linear system (see Jacobsen et al. (2023)), giving pairs (\mathbf{K}, \mathbf{p}) with $\mathbf{K}, \mathbf{p} \in \mathbb{R}^{n \times n}$. We consider a U-Net (Ronneberger et al., 2015) architecture with 64×64 pixels as in- and output dimensions that align with the considered grid and allow for the same residual evaluation via finite differences also used to create the dataset. We increase the variance of the virtual residual likelihood by setting $c = 10^{-3}$ for the mean estimation and $c = 10^{-5}$ for the sample estimation. Throughout the experiments, we observed that the mean estimation performs best with a variance of the residual likelihood around two orders of magnitude smaller than the best performance for the sample estimation. Further details of the considered setup and implementation are given in Appendices A.6.2 and A.7.

Results. To evaluate the performance of our proposed PIDM, we benchmark it against three relevant setups: (i) a model trained on the standard (purely data-driven) objective equation 4, (ii) a “physics-guided” model trained on the standard objective but using residual information as guidance, as proposed by Shu et al. (2023), and (iii) a model similar to (i) but with first-order residual corrections during inference, as described in CoCoGen (Jacobsen et al., 2023) (see Appendix A.6.2 for further details). We examine the performance of the different diffusion model variants by tracking the evolution of the residual error of generated samples from the learned distributions $\mathbf{x}_0 \sim p_{\theta}(\mathbf{x}_0)$ alongside the test data loss throughout the training process in Figure 2. The PIDM showcases a remarkable improvement, reducing the residual error by around two orders of magnitude in comparison to the standard diffusion model. This holds for both the mean and sample (DDIM) estimation, where the former outperforms the latter variant in terms of the residual error. The mean estimation takes around 23% longer to train compared to the standard diffusion model due to the residual evaluation, whereas the sample estimation took around 69% longer due to the additional forward pass (see Appendix A.4). We could not observe a significant reduction of the residual error for the physics-guided model (Shu et al., 2023). We hypothesize that the mere inclusion of gradient information into the model seems insufficient because it does not truly enforce residual minimization. CoCoGen (Jacobsen et al., 2023) shows a moderate improvement but does not bring the residual close to values where the samples could be considered physically consistent. We display generated samples in Figure 3 and Appendix A.8.1 for both the standard diffusion model and PIDM, confirming a drastically reduced residual error over the whole domain. Figures 2b, 3b and 8 confirm that the


324
325
326
327
328
329
330
331
332
333
334
335
336
337
338



339
340
341
342
343
344
345
346
347
348

Figure 2: Evaluation of the residual error (a) and test data loss (b) during training. In (a), we generate 16 samples every 10k training iterations and plot the average (solid lines) and individual (dots) residual errors for the standard diffusion model (‘Diffusion’), the physics-guided model (‘PG-Diffusion’) (Shu et al., 2023), CoCoGen (Jacobsen et al., 2023), and the proposed PIDM using either mean or sample estimation (‘PIDM-ME’ and ‘PIDM-SE’, respectively). Note that for CoCoGen not all samples converged, so that we excluded the non-converged data from the indicated average. In (b), we plot the data loss evaluated on a test set for the proposed PIDM variants and those frameworks that differ from ours during training.

349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368



369
370
371
372
373
374
375
376
377

Figure 3: Generated permeability and pressure fields as well as the corresponding residual error from diffusion models trained on the Darcy flow dataset, where (a) is sampled from a standard diffusion model and (b) from our proposed PIDM with mean estimation. Additional samples are shown in Appendix A.8.1.

PIDMs maintain the generative diversity with both estimation techniques and do not collapse to a single solution instance. It successfully generates permeability (and pressure) fields that adhere to the data distribution (besides respecting the PDE).

For any multi-objective loss as given in equation 13, a trade-off between the data and residual minimization is generally expected, depending on their relative weighting dictated by c (?). Interestingly, Figure 2b shows that the PIDM with mean estimation eventually recovers a similar test data loss as the vanilla diffusion model and is significantly less prone to overfitting. Even more notable is that the PIDM with sample estimation initially has a similar test data loss trajectory as the vanilla model—despite the additional loss term—and also remains more robust to overfitting. This is strong evidence of the consistent data and physics loss (see Appendix A.1). Thus, sample estimation drastically improves the alignment of the two losses with only a single additional forward pass. The increased robustness to overfitting in both mean and sample estimation is a clear indicator that the model learns a more robust internal representation of the data distribution, as it is forced to generate samples that adhere to the true distribution or, equivalently, the underlying data generation mechanism—viz. the known constraints that govern the system. Thus, adding the physical (residual) loss may benefit the generative performance by enhancing its generalization capability.

4.2 TOPOLOGY OPTIMIZATION

Setup. We consider 2D structural topology optimization (TO) as a second example. In this setting, the goal is to find the optimal material distribution $\rho_{\text{opt}}(\xi)$ that maximizes the mechanical stiffness, or equivalently, minimizes the compliance of a structure under a set of constraints that typically consist of mechanical equilibrium, boundary conditions, and a volume constraint. This is classically solved via the SIMP method based on a finite element (FE) discretization (?). We again consider a square domain $\Omega = [0, 1]^2$ and benchmark our proposed PIDM to state-of-the-art frameworks (Mazé & Ahmed, 2023; Giannone et al., 2023) that also provide a dataset consisting of 30,000 optimized structures with various boundary conditions and volume constraints and two proposed test scenarios with in- and out-of-distribution boundary conditions. We train a U-Net architecture similar to the one in Section 4.1 but with a larger latent dimension and additional in- and output channels on this dataset. To ensure consistent evaluation of the residuals, we do *not* use finite differences but interpret the pixels as nodes of the underlying FE mesh to assemble a consistent stiffness matrix equivalent to the one used to generate the data. We scale the variance of the residual likelihood with $c = 0.01$ and introduce the volume constraint as an additional equality constraint with $c = 0.1$ (as the optimal topology will contain the maximum allowed material). We also introduce a slight bias to minimize compliance (setting $\lambda = 10^{-6}$ for the optimization objective, see Appendix A.5.2). Further details of the setup, model architecture, residual evaluation, and implementation are presented in Appendices A.6.3 and A.7.

Results. We evaluate the performance of our proposed PIDM in comparison to the standard diffusion model, PG-Diffusion (Shu et al., 2023) and CoCoGen (Jacobsen et al., 2023), as well as two recent variants specifically tailored to topology optimization. These propose to modify the sampling process by either using additional guidance models to reduce compliance and improve manufacturability (Mazé & Ahmed, 2023) or enforcing the denoising trajectory to be closer to an iterative optimization trajectory (Giannone et al., 2023). Both methods require auxiliary datasets which complicate model training—in contrast, our method is a much simpler and well-motivated extension of the standard training without requiring any additional data or models. While training of the ‘main’ diffusion model takes longer for the PIDM due to the additional computational complexity and memory requirements, we note that inference is identical to the standard diffusion model and does not require additional surrogate models as, e.g., in Mazé & Ahmed (2023). Results are shown in Table 1 for the PIDM with sample estimation, which we observed to outperform the mean estimation slightly. We evaluate relevant performance metrics on the two test sets (with seen and unseen boundary conditions, respectively), with generated samples presented in Figure 4 and Appendix A.8.2. Importantly, the PIDM provides not only the optimized designs but also the displacement fields, which are of significant interest for mechanical analysis, e.g., to estimate the stress distribution in the structure. We observe a residual reduced by 33% and 35% compared to the standard diffusion model on the in- and out-of-distribution set, respectively, highlighting closer adherence to the mechanical equilibrium. Somewhat surprisingly, we observe that the standard model performs very well in terms of the optimization objective, but also emphasize that it has an increased volume fraction error—thus results are not fully comparable, as more volume will allow for stiffer structures. Contrary, the PIDM only has a slight increase in the volume fraction error but significantly outperforms previous frameworks (Mazé & Ahmed, 2023; Giannone et al., 2023) in terms of compliance minimization, despite the absence of auxiliary data and surrogate models. Lastly, adaptations of PG-Diffusion (Shu et al., 2023)

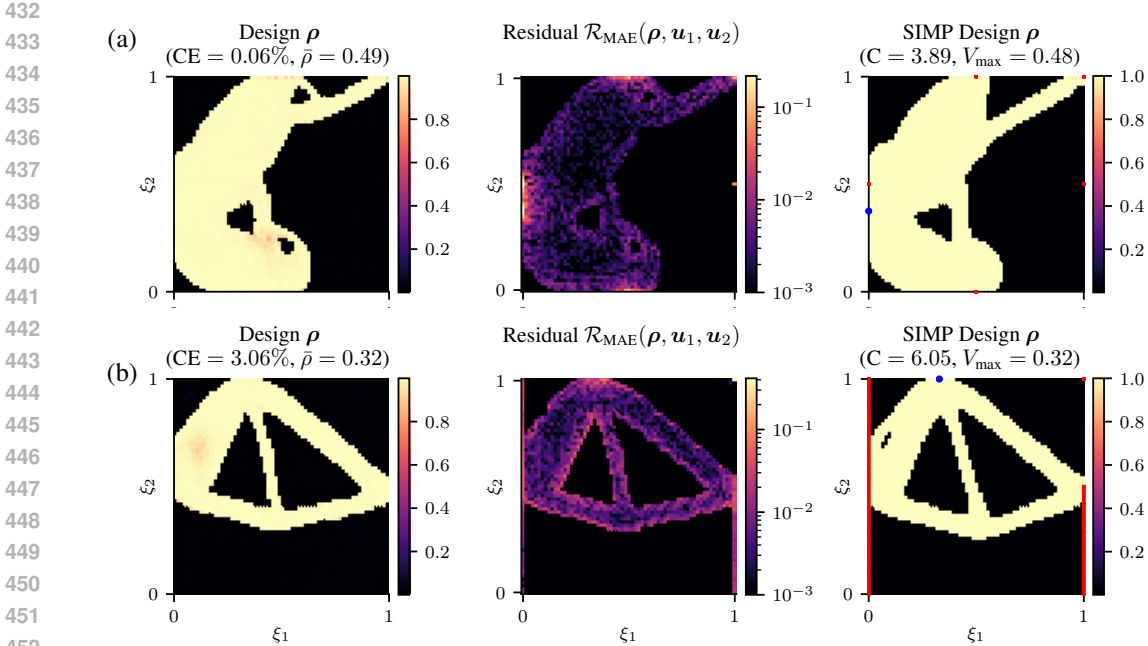


Figure 4: Generated designs, including the compliance error CE and volume $\bar{\rho}$, and the residual error (based on the displacement fields, not shown) from diffusion models trained on the SIMP dataset and the corresponding SIMP design, including the compliance C and volume V_{\max} . We plot a sample (a) from a standard diffusion model and (b) from our proposed PIDM. All samples are conditioned on the out-of-distribution test set. In the SIMP design, we indicate the applied load by a blue dot, and the given boundary conditions in red.

Table 1: Performance comparison of diffusion model variants for topology optimization. We consider in- and out-of-distribution boundary conditions as described in Mazé & Ahmed (2023) and provide the median \mathcal{R}_{MAE} of the predicted solution fields (where applicable), the median compliance error (MDN % CE) and the mean volume fraction error (% VFE). *Giannone et al. (2023) further improve their model by running additional SIMP post-processing steps, but for a consistent comparison, we consider unprocessed samples from the model.

Model	Size	in-distribution			out-of-distribution		
		$\mathcal{R}_{\text{MAE}} \downarrow$	MDN % CE \downarrow	% VFE \downarrow	$\mathcal{R}_{\text{MAE}} \downarrow$	MDN % CE \downarrow	% VFE \downarrow
Diffusion	136M	1.86e-3	-0.2	2.93	1.97e-3	0.3	2.80
PG-Diffusion (Shu et al., 2023)	136M	1.82e-3	0.09	3.59	1.92e-3	0.81	3.23
CoCoGen (Jacobsen et al., 2023)	136M	1.51e-3	0.14	4.00	1.56e-3	0.58	3.64
TopoDiff-G (Mazé & Ahmed, 2023)	239M	-	0.83	1.49	-	1.82	1.80
DOM* (Giannone et al., 2023)	121M	-	0.74	1.52	-	3.47	1.59
PIDM (ours)	136M	1.24e-3	0.06	2.25	1.29e-3	0.56	1.91

and CoCoGen (Jacobsen et al., 2023) also cannot match the PIDM in terms of the residual (on both test sets), though CoCoGen can reduce it slightly. Additionally, the median compliance and especially the volume fraction error perform significantly worse.

5 RELATED WORK

Conceptually closest to our work are two recent contributions by Shu et al. (2023) and Jacobsen et al. (2023), to which we compare the proposed PIDM extensively in Section 4. Our studies indicate that the PIDM significantly reduces the residual error of generated samples compared to both variants, especially for the Darcy flow study. Focusing on topology optimization, two recent contributions (Mazé & Ahmed, 2023; Giannone et al., 2023) propose improvements that require auxiliary data

and/or surrogate models, while we here show that the PIDM, solely by being *physics-informed*, significantly outperforms both methods in terms of the optimal stiffness. In a broader context, Wang et al. (2023) optimized the conditioning variable to create an online dataset of soft robot designs, which includes physical performance and may generate samples with improved physical utility. Yet, the focus was on leveraging pre-trained models to create diverse 3D shapes. Yuan et al. (2022) leveraged diffusion models for human motion synthesis, where inference is altered by projecting the intermediate steps to a physically plausible motion that is verified via a reinforcement learning approach. Similarly, Christopher et al. (2024) projected intermediate steps to the closest point in a feasible set, which is generally unknown for more complicated constraints such as PDEs. In general, such post-processing methods may indeed mitigate some of the mismatches of the generated samples (or fulfill them exactly if the constraints are sufficiently simple). Yet, they are fundamentally limited, as they do not address the underlying distribution learned by the model.

6 CONCLUSION

We have unified the data-driven perspective of diffusion models with a physics-informed paradigm, enabling the models to internalize the constraints that generated samples must adhere to. Our framework significantly outperforms purely data-driven models and prior work, as verified by two highly relevant case studies, and numerical evidence hints that the PIDM obtains a more robust representation of the data distribution that is less prone to overfitting. We hope this work stimulates others to extend their generative model training objective when, besides data, further information—be it in the form of PDEs or other constraints—on the generated samples is available, as is often the case within the realm of scientific machine learning. Future work may explore more sophisticated virtual likelihood variance schedulers, which we here simply coupled to the standard denoising process with a scaling estimated by a parameter sweep. More generally, architectures with a consistent residual evaluation that do not operate on a fixed regular grid should be considered, currently restricting the studies to geometrically simple domains. Remedy can be found in graph-based architectures that can handle arbitrary meshes (Gao et al., 2022) or by employing implicit encodings of coordinates (Liu et al., 2023). Additionally, coordinate-based representations allow for exact calculation of the gradients required to evaluate the imposed PDEs via automatic differentiation, though at increased computational cost.

REFERENCES

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer Normalization. 2016. URL <http://arxiv.org/abs/1607.06450>.
- Matthias Baer. {findiff} Software Package, 2018. URL <https://github.com/maroba/findiff>.
- Jan-Hendrik Bastek and Dennis M. Kochmann. Inverse design of nonlinear mechanical metamaterials via video denoising diffusion models. *Nature Machine Intelligence*, 5(12):1466–1475, 2023. ISSN 2522-5839. doi: 10.1038/s42256-023-00762-x. URL <https://www.nature.com/articles/s42256-023-00762-x>.
- Julius Berner, Lorenz Richter, and Karen Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=oYIjw37pTP>.
- Markus J. Buehler. Modeling Atomistic Dynamic Fracture Mechanisms Using a Progressive Transformer Diffusion Model. *Journal of Applied Mechanics*, 89(12), 2022. ISSN 0021-8936. doi: 10.1115/1.4055730. URL <https://asmedigitalcollection.asme.org/appliedmechanics/article/89/12/121009/1146377/Modeling-Atomistic-Dynamic-Fracture-Mechanisms>.
- Jacob K Christopher, Stephen Baek, and Ferdinando Fioretto. Projected Generative Diffusion Models for Constraint Satisfaction. 2024. URL <http://arxiv.org/abs/2402.03559>.
- Wojciech Marian Czarnecki, Simon Osindero, Max Jaderberg, Grzegorz Świrszcz, and Razvan Pascanu. Sobolev Training for Neural Networks. 2017. URL <http://arxiv.org/abs/1706.04859>.

- 540 Prafulla Dhariwal and Alex Nichol. Diffusion Models Beat GANs on Image Synthesis. 2021. URL
541 <http://arxiv.org/abs/2105.05233><http://arxiv.org/abs/2010.02502>.
- 542
- 543 Christian Dürer, Paul Seibert, Dennis Rücker, Stephanie Handford, Markus Kästner, and Maik Gude.
544 Conditional diffusion-based microstructure reconstruction. *Materials Today Communications*,
545 35:105608, 2023. ISSN 23524928. doi: 10.1016/j.mtcomm.2023.105608. URL [https://](https://linkinghub.elsevier.com/retrieve/pii/S2352492823002982)
546 linkinghub.elsevier.com/retrieve/pii/S2352492823002982.
- 547 Stefan Elfving, Eiji Uchibe, and Kenji Doya. Sigmoid-Weighted Linear Units for Neural Network
548 Function Approximation in Reinforcement Learning. 2017. URL [http://arxiv.org/abs/](http://arxiv.org/abs/1702.03118)
549 [1702.03118](http://arxiv.org/abs/1702.03118).
- 550 Han Gao, Matthew J. Zahr, and Jian Xun Wang. Physics-informed graph neural Galerkin networks:
551 A unified framework for solving PDE-governed forward and inverse problems. *Computer Methods*
552 *in Applied Mechanics and Engineering*, 390:114502, 2022. ISSN 00457825. doi: 10.1016/j.cma.
553 2021.114502. URL <https://doi.org/10.1016/j.cma.2021.114502>.
- 554
- 555 Xiang Gao, Meera Sitharam, and Adrian E. Roitberg. Bounds on the Jensen Gap, and Implications for
556 Mean-Concentrated Distributions. dec 2017. URL <http://arxiv.org/abs/1712.05267>.
- 557 Giorgio Giannone, Akash Srivastava, Ole Winther, and Faez Ahmed. Aligning Optimization
558 Trajectories with Diffusion Models for Constrained Design Generation. In A Oh, T Neu-
559 mann, A Globerson, K Saenko, M Hardt, and S Levine (eds.), *Advances in Neural Informa-*
560 *tion Processing Systems*, volume 36, pp. 51830–51861. Curran Associates, Inc., 2023.
561 URL [https://proceedings.neurips.cc/paper/_files/paper/2023/file/](https://proceedings.neurips.cc/paper/_files/paper/2023/file/a2fe4bb50fc6f3564cee1551d6309fea-Paper-Conference.pdf)
562 [a2fe4bb50fc6f3564cee1551d6309fea-Paper-Conference.pdf](https://proceedings.neurips.cc/paper/_files/paper/2023/file/a2fe4bb50fc6f3564cee1551d6309fea-Paper-Conference.pdf).
- 563 Juan Gómez and Nicolás Guarín-Zapata. SolidsPy: 2D-Finite Element Analysis with Python, 2018.
564 URL <https://github.com/AppliedMechanics-EAFIT/SolidsPy>.
- 565
- 566 Tiankai Hang, Shuyang Gu, Chen Li, Jianmin Bao, Dong Chen, Han Hu, Xin Geng, and Baining Guo.
567 Efficient Diffusion Training via Min-SNR Weighting Strategy. 2023. URL [http://arxiv.](http://arxiv.org/abs/2303.09556)
568 [org/abs/2303.09556](http://arxiv.org/abs/2303.09556).
- 569 Jonathan Ho and Tim Salimans. Classifier-Free Diffusion Guidance. 2022. URL [http://arxiv.](http://arxiv.org/abs/2207.12598)
570 [org/abs/2207.12598](http://arxiv.org/abs/2207.12598).
- 571
- 572 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. *Advances*
573 *in Neural Information Processing Systems*, (NeurIPS 2020):1–25, 2020. ISSN 10495258. URL
574 <http://arxiv.org/abs/2006.11239><http://arxiv.org/abs/2010.02502>.
- 575 Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P.
576 Kingma, Ben Poole, Mohammad Norouzi, David J. Fleet, and Tim Salimans. Imagen Video:
577 High Definition Video Generation with Diffusion Models. pp. 90–91, 2022a. URL [http://](http://arxiv.org/abs/2210.02303)
578 arxiv.org/abs/2210.02303.
- 579 Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J.
580 Fleet. Video Diffusion Models. 2022b. URL <http://arxiv.org/abs/1803.02155>[http://](http://arxiv.org/abs/2204.03458)
581 arxiv.org/abs/2204.03458.
- 582
- 583 Emiel Hoogeboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant Diffusion
584 for Molecule Generation in 3D. 2022. URL <http://arxiv.org/abs/2203.17003>.
- 585 Christian Jacobsen, Yilin Zhuang, and Karthik Duraisamy. CoCoGen: Physically-Consistent and
586 Conditioned Score-based Generative Models for Forward and Inverse Problems. pp. 1–25, 2023.
587 URL <http://arxiv.org/abs/2312.10527>.
- 588
- 589 Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are
590 RNNs: Fast Autoregressive Transformers with Linear Attention. 2020. URL [http://arxiv.](http://arxiv.org/abs/2006.16236)
591 [org/abs/2006.16236](http://arxiv.org/abs/2006.16236).
- 592 Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *3rd International*
593 *Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, pp. 1–15,
2014. URL <http://arxiv.org/abs/1412.6980>.

- 594 Diederik P Kingma and Max Welling. Auto-Encoding Variational Bayes. 2013. URL <http://arxiv.org/abs/1312.6114>.
595
596
- 597 Diederik P. Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational Diffusion Models.
598 (NeurIPS), jul 2021. URL <http://arxiv.org/abs/2107.00630>.
599
- 600 Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. DiffWave: A Versatile
601 Diffusion Model for Audio Synthesis. pp. 1–17, 2020. URL <http://arxiv.org/abs/2009.09761>.
602
- 603 Tianyi Li, Alessandra S. Lanotte, Michele Buzzicotti, Fabio Bonaccorso, and Luca Biferale. Multi-
604 Scale Reconstruction of Turbulent Rotating Flows with Generative Diffusion Models. *Atmosphere*,
605 15(1):60, 2023. ISSN 2073-4433. doi: 10.3390/atmos15010060. URL <https://www.mdpi.com/2073-4433/15/1/60>.
606
- 607 Marten Lienen, David Lüdke, Jan Hansen-Palmus, and Stephan Günnemann. From Zero to Turbu-
608 lence: Generative Modeling for 3D Flow Simulation. 2023. URL <http://arxiv.org/abs/2306.01776>.
609
610
- 611 Hongying Liu, Zekun Li, Fanhua Shang, Yuanyuan Liu, Liang Wan, Wei Feng, and
612 Radu Timofte. Arbitrary-scale super-resolution via deep learning: A comprehensive survey.
613 *Information Fusion*, pp. 102015, 2023. ISSN 15662535. doi: 10.1016/j.inffus.2023.102015
614 <https://doi.org/10.1016/j.inffus.2023.102015>
615 <https://linkinghub.elsevier.com/retrieve/pii/S1566253523003317>.
616
- 617 François Mazé and Faez Ahmed. Diffusion Models Beat GANs on Topology Optimization. *Pro-*
618 *ceedings of the 37th AAAI Conference on Artificial Intelligence, AAAI 2023*, 37:9108–9116, 2023.
619 ISSN 2159-5399. doi: 10.1609/aaai.v37i8.26093.
- 620 Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya
621 Sutskever, and Mark Chen. GLIDE: Towards Photorealistic Image Generation and Editing with
622 Text-Guided Diffusion Models. 2021. URL <http://arxiv.org/abs/2112.10741>.
623
- 624 Zhenguo Nie, Tong Lin, Haoliang Jiang, and Levent Burak Kara. TopologyGAN: Topol-
625 ogy Optimization Using Generative Adversarial Networks Based on Physical Fields
626 Over the Initial Domain. *Journal of Mechanical Design*, 143(3), mar 2021. ISSN
627 1050-0472. doi: 10.1115/1.4049533. URL [https://asmedigitalcollection.
628 asme.org/mechanicaldesign/article/doi/10.1115/1.4049533/1094063/
629 TopologyGAN-Topology-Optimization-Using-Generative](https://asmedigitalcollection.asme.org/mechanicaldesign/article/doi/10.1115/1.4049533/1094063/TopologyGAN-Topology-Optimization-Using-Generative).
630
- 631 Chenhao Niu, Yang Song, Jiaming Song, Shengjia Zhao, Aditya Grover, and Stefano Ermon. Per-
632 mutation Invariant Graph Generation via Score-Based Generative Modeling. 108, 2020. URL
633 <http://arxiv.org/abs/2003.00638>.
634
- 635 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
636 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward
637 Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang,
638 Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning
639 library. *Advances in Neural Information Processing Systems*, 2019. ISSN 23318422. URL
640 <http://arxiv.org/abs/1912.01703>.
641
- 642 M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning
643 framework for solving forward and inverse problems involving nonlinear partial differential
644 equations. *Journal of Computational Physics*, 378:686–707, 2019. ISSN 00219991. doi: 10.
645 1016/j.jcp.2018.10.045. URL [https://linkinghub.elsevier.com/retrieve/pii/
646 S0021999118307125](https://linkinghub.elsevier.com/retrieve/pii/S0021999118307125).
647
- 648 Maximilian Rixner and Phaedon Stelios Koutsourelakis. A probabilistic generative model for semi-
649 supervised training of coarse-grained surrogates and enforcing physical constraints through virtual
650 observables. *Journal of Computational Physics*, 434, 2021. ISSN 10902716. doi: 10.1016/j.jcp.
651 2021.110218.

- 648 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
649 Resolution Image Synthesis with Latent Diffusion Models. 2021. URL [http://arxiv.org/
650 abs/2112.10752](http://arxiv.org/abs/2112.10752).
651
- 652 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical
653 Image Segmentation. 2015. URL <http://arxiv.org/abs/1505.04597>.
654
- 655 Tim Salimans and Jonathan Ho. Progressive Distillation for Fast Sampling of Diffusion Models. pp.
656 1–21, feb 2022. URL <http://arxiv.org/abs/2202.00512>.
657
- 658 Mohammed Sardar, Alex Skillen, Małgorzata J. Zimoń, Samuel Draycott, and Alistair Revell.
659 Spectrally Decomposed Diffusion Models for Generative Turbulence Recovery. 2023. URL
<http://arxiv.org/abs/2312.15029>.
660
- 661 Dule Shu, Zijie Li, and Amir Barati Farimani. A physics-informed diffusion model for high-
662 fidelity flow field reconstruction. *Journal of Computational Physics*, 478:111972, 2023. ISSN
663 10902716. doi: 10.1016/j.jcp.2023.111972. URL [https://doi.org/10.1016/j.jcp.
664 2023.111972](https://doi.org/10.1016/j.jcp.2023.111972).
665
- 666 Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep Unsupervised
667 Learning using Nonequilibrium Thermodynamics. *ICML*, 2015. URL [http://arxiv.org/
668 abs/1503.03585](http://arxiv.org/abs/1503.03585).
669
- 670 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising Diffusion Implicit Models. oct 2020.
671 URL <http://arxiv.org/abs/2010.02502>.
672
- 673 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and
674 Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations. In
International Conference on Learning Representations, 2021. URL [https://openreview.
675 net/forum?id=PXTIG12RRHS](https://openreview.net/forum?id=PXTIG12RRHS).
676
- 677 Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency Models. mar 2023. URL
678 <http://arxiv.org/abs/2303.01469>.
679
- 680 Francisco Vargas, Will Sussman Grathwohl, and Arnaud Doucet. Denoising Diffusion Samplers.
681 In *The Eleventh International Conference on Learning Representations*, 2023. URL [https://
682 openreview.net/forum?id=8pvnfTAbulf](https://openreview.net/forum?id=8pvnfTAbulf).
683
- 684 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz
685 Kaiser, and Illia Polosukhin. Attention Is All You Need. 2017. URL [http://arxiv.org/
686 abs/1706.03762](http://arxiv.org/abs/1706.03762).
687
- 688 Pascal Vincent. A Connection Between Score Matching and Denoising Autoencoders. *Neural
689 Computation*, 23(7):1661–1674, jul 2011. ISSN 0899-7667. doi: 10.1162/NECO_a_00142. URL
690 <https://direct.mit.edu/neco/article/23/7/1661-1674/7677>.
691
- 692 Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau,
693 Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt,
694 Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric
695 Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas,
696 Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris,
697 Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, Aditya Vijaykumar,
698 Alessandro Pietro Bardelli, Alex Rothberg, Andreas Hilboll, Andreas Kloeckner, Anthony Scopatz,
699 Antony Lee, Ariel Rokem, C. Nathan Woods, Chad Fulton, Charles Masson, Christian Häggström,
700 Clark Fitzgerald, David A. Nicholson, David R. Hagen, Dmitrii V. Pasechnik, Emanuele Olivetti,
701 Eric Martin, Eric Wieser, Fabrice Silva, Felix Lenders, Florian Wilhelm, G. Young, Gavin A.
Price, Gert-Ludwig Ingold, Gregory E. Allen, Gregory R. Lee, Hervé Audren, Irvin Probst,
Jörg P. Dietrich, Jacob Silterra, James T Webber, Janko Slavič, Joel Nothman, Johannes Buchner,
Johannes Kulick, Johannes L. Schönberger, José Vinícius de Miranda Cardoso, Joscha Reimer,
Joseph Harrington, Juan Luis Cano Rodríguez, Juan Nunez-Iglesias, Justin Kuczynski, Kevin
Tritz, Martin Thoma, Matthew Newville, Matthias Kümmerer, Maximilian Bolingbroke, Michael
Tartre, Mikhail Pak, Nathaniel J. Smith, Nikolai Nowaczyk, Nikolay Shebanov, Oleksandr Pavlyk,

- 702 Per A. Brodtkorb, Perry Lee, Robert T. McGibbon, Roman Feldbauer, Sam Lewis, Sam Tygier,
703 Scott Sievert, Sebastiano Vigna, Stefan Peterson, Surhud More, Tadeusz Pudlik, Takuya Oshima,
704 Thomas J. Pingel, Thomas P. Robitaille, Thomas Spura, Thouis R. Jones, Tim Cera, Tim Leslie,
705 Tiziano Zito, Tom Krauss, Utkarsh Upadhyay, Yaroslav O. Halchenko, and Yoshiki Vázquez-
706 Baeza. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*,
707 17(3):261–272, 2020. ISSN 1548-7091. doi: 10.1038/s41592-019-0686-2. URL <https://www.nature.com/articles/s41592-019-0686-2>.
- 709 Nikolaos N. Vlassis and WaiChing Sun. Denoising diffusion algorithm for inverse design of mi-
710 crostructures with fine-tuned nonlinear material properties. *Computer Methods in Applied Mechan-*
711 *ics and Engineering*, 413:116126, 2023. ISSN 00457825. doi: 10.1016/j.cma.2023.116126. URL
712 <https://linkinghub.elsevier.com/retrieve/pii/S0045782523002505>.
- 713 Phil Wang. Implementation of Imagen, Google’s Text-to-Image Neural Network that beats DALL-E2,
714 in Pytorch, 2022. URL <https://github.com/lucidrains/imagen-pytorch>.
- 716 Tsun-hsuan Wang, Juntian Zheng, Pingchuan Ma, Yilun Du, Byungchul Kim, Andrew Spielberg,
717 Joshua Tenenbaum, Chuang Gan, and Daniela Rus. DiffuseBot: Breeding Soft Robots With
718 Physics-Augmented Generative Diffusion Models. (1830901):1–26, 2023. URL <http://arxiv.org/abs/2311.17053>.
- 720 Tong Wu, Zhihao Fan, Xiao Liu, Yeyun Gong, Yelong Shen, Jian Jiao, Hai-Tao Zheng, Juntao Li,
721 Zhongyu Wei, Jian Guo, Nan Duan, and Weizhu Chen. AR-Diffusion: Auto-Regressive Diffusion
722 Model for Text Generation. 2023. URL <http://arxiv.org/abs/2305.09515>.
- 723 Yuxin Wu and Kaiming He. Group Normalization. 2018. URL <http://arxiv.org/abs/1803.08494>.
- 726 Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi Jaakkola. Crystal Diffusion
727 Variational Autoencoder for Periodic Material Generation. pp. 1–20, 2021. URL <http://arxiv.org/abs/2110.06197>.
- 729 Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. GeoDiff: a
730 Geometric Diffusion Model for Molecular Conformation Generation. (2021):1–19, 2022. URL
731 <http://arxiv.org/abs/2203.02923>.
- 733 Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang,
734 Bin Cui, and Ming-Hsuan Yang. Diffusion Models: A Comprehensive Survey of Methods and
735 Applications. *ACM Computing Surveys*, 56(4):1–39, 2024. ISSN 0360-0300. doi: 10.1145/3626235.
736 URL <https://dl.acm.org/doi/10.1145/3626235>.
- 737 Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. PhysDiff: Physics-Guided Human
738 Motion Diffusion Model. 2022. URL <http://arxiv.org/abs/2212.02500>.
- 739 Qinsheng Zhang and Yongxin Chen. Fast Sampling of Diffusion Models with Exponential Integrator.
740 apr 2022a. URL <http://arxiv.org/abs/2204.13902>.
- 742 Qinsheng Zhang and Yongxin Chen. Path Integral Sampler: A Stochastic Control Approach For
743 Sampling. In *International Conference on Learning Representations*, 2022b. URL https://openreview.net/forum?id={_}uCb2ynRu7Y.
- 745 Yin hao Zhu and Nicholas Zabaras. Bayesian deep convolutional encoder–decoder networks for
746 surrogate modeling and uncertainty quantification. *Journal of Computational Physics*, 366:415–
747 447, 2018. ISSN 00219991. doi: 10.1016/j.jcp.2018.04.018. URL <https://linkinghub.elsevier.com/retrieve/pii/S0021999118302341>.

750 A APPENDIX

753 A.1 CONSISTENCY

754 We here give a theoretical argument for the fact that, in an ideal setting, diffusion models trained
755 via score-matching with an additional virtual (residual) likelihood term continue to recover the true

distribution. Adopting the perspective of score-based models (Song et al., 2021), we consider a continuous time variable $t \in [0, T]$ instead of the discrete setting. The forward diffusion process of some data $\mathbf{x} \in \mathbb{R}^D$ can then generally be described by the stochastic differential equation (SDE)

$$d\mathbf{x} = \mathbf{F}_t \mathbf{x} dt + \mathbf{G}_t d\mathbf{w}, \quad (14)$$

where \mathbf{w} is a Wiener process, and the drift and diffusion coefficients $\mathbf{F}_t \in \mathbb{R}^{D \times D}$ and $\mathbf{G}_t \in \mathbb{R}^{D \times D}$, respectively, are chosen so that the transition kernel is a simple Gaussian that can be computed in closed form. Here, the popular diffusion variants proposed by Ho et al. (2020) and DDIM (Song et al., 2021) correspond to their continuous counterparts $\mathbf{F}_t = \frac{1}{2} \frac{d \log \alpha_t}{dt} \mathbf{I}$ and $\mathbf{G}_t = \sqrt{-\frac{d \log \alpha_t}{dt}} \mathbf{I}$, where α_t is continuously decreasing from $\alpha_0 \approx 1$ to $\alpha_T \approx 0$. There is no unique reverse diffusion process, but common choices of the corresponding reverse SDE are included in the parameterization proposed by Zhang & Chen (2022a),

$$d\mathbf{x} = \left[\mathbf{F}_t \mathbf{x} - \frac{1 + \lambda^2}{2} \mathbf{G}_t \mathbf{G}_t^T \nabla \log p_t(\mathbf{x}) \right] d\bar{t} + \lambda \mathbf{G}_t d\bar{\mathbf{w}}, \quad (15)$$

where $\bar{\mathbf{w}}$ is a Wiener process that runs, as $d\bar{t}$, backwards in time, and $\lambda \geq 0$. If the score $\nabla \log p_t(\mathbf{x})$ is known, we can generate new samples by sampling from the known prior distribution $p(\mathbf{x}_T)$ and applying equation 15. DDPM (Ho et al., 2020) and (deterministic) DDIM (Song et al., 2020) then correspond to certain parameterizations ($\lambda = 1$ and $\lambda = 0$, respectively) and discretizations of equation 15 (Song et al., 2021; Zhang & Chen, 2022a). We can show that a straightforward extension of the usual score-matching objective with the virtual residual likelihood is consistent as it continues to recover the data distribution:

Proposition 1. (Consistency) *Let $p(\mathbf{x}_0)$ be a distribution with samples $\mathbf{x}_0 \sim p(\mathbf{x}_0)$ satisfying some constraint $\mathcal{R}(\mathbf{x}_0) = \mathbf{0}$. Consider*

$$\mathbf{s}_{\text{opt}} = \arg \min_{\mathbf{s}} \mathbb{E}_{t \sim \text{Unif}[0, T]} \mathbb{E}_{p(\mathbf{x}_0) p(\mathbf{x}_t | \mathbf{x}_0)} \left[\Lambda(t) \|\nabla \log p(\mathbf{x}_t | \mathbf{x}_0) - \mathbf{s}(\mathbf{x}_t, t)\|^2 - \log q_{\mathcal{R}}(\hat{\mathbf{r}} | \mathbf{x}_0^*(\mathbf{x}_t)) \right], \quad (16)$$

where $\Lambda(t) > 0$ is a time-dependent weight and \mathbf{x}_0^* is obtained by solving the reverse SDE equation 15 initiated at \mathbf{x}_t with score $\mathbf{s}(\mathbf{x}_t, t)$. Then, solving the reverse SDE equation 15 from $\mathbf{x}_T \sim p(\mathbf{x}_T)$ with \mathbf{s}_{opt} as the score for \mathbf{x}_0 corresponds to sampling from $p(\mathbf{x}_0)$.

Proof. It is well-known that score matching $\nabla \log p(\mathbf{x}_t | \mathbf{x}_0)$ is equivalent to matching $\nabla \log p(\mathbf{x}_t)$ (Vincent, 2011). Assuming perfect recovery of the score, i.e., $\mathbf{s}(\mathbf{x}, t) = \nabla \log p(\mathbf{x}_t)$ for all \mathbf{x}_t, t , the marginal distribution $p^*(\mathbf{x}_t)$ of equation 15 matches the forward diffusion $p(\mathbf{x}_t)$ for all $0 \leq t \leq T$ independent of λ when $p(\mathbf{x}_T) = p^*(\mathbf{x}_T)$ (Zhang & Chen, 2022a, Prop. 1). To generate new samples, we may start from any latent $\mathbf{x}_t^* \sim p^*(\mathbf{x}_t)$, which can hence equivalently be obtained via the forward marginal $\mathbf{x}_t \sim p(\mathbf{x}_t)$, and solve equation 15 for $\mathbf{x}_0^* \sim p(\mathbf{x}_0)$. As we consider virtual observables $\hat{\mathbf{r}} = \mathbf{0}$ introduced via $q_{\mathcal{R}}(\hat{\mathbf{r}} | \mathbf{x}_0^*) = \mathcal{N}(\hat{\mathbf{r}}; \mathcal{R}(\mathbf{x}_0^*), \sigma^2 \mathbf{I})$ we have

$$-\log q_{\mathcal{R}}(\hat{\mathbf{r}} | \mathbf{x}_0^*) = \frac{1}{2} \|\mathcal{R}(\mathbf{x}_0^*) / \sigma^2\|^2 + C, \quad (17)$$

where C is a constant that does not depend on \mathbf{x}_0^* (and hence \mathbf{s}). Since $\mathbf{x}_0^* \sim p(\mathbf{x}_0)$ satisfies $\mathcal{R}(\mathbf{x}_0^*) = \mathbf{0}$ by assumption, the optimal score model is also a minimizer of the (negative) virtual log-likelihood equation 17 and \mathbf{s}_{opt} recovers the optimal score. Hence, solving equation 15 from $\mathbf{x}_T \sim p(\mathbf{x}_T)$ for \mathbf{x}_0 with \mathbf{s}_{opt} as the score generates samples from the true distribution $\mathbf{x}_0 \sim p(\mathbf{x}_0)$. \square

Remark. In practice, equation 15 has to be discretized, since no closed-form solution is available. Hence, even for a perfect score model, this numerical approximation will introduce a bias (Zhang & Chen, 2022a). Also, note that we are free to choose DDPM (Ho et al., 2020) or DDIM (Song et al., 2020) to discretize equation 15 since the optimal score is unaffected by this choice.

Lastly, we emphasize that minimizing the residual error on some intermediate latent variables introduces inconsistencies. Assume some general forward conditional marginal of the form $p(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \alpha_t \mathbf{x}_0, \sigma_t^2 \mathbf{I})$, from which we can sample via $\mathbf{x}_t = \alpha_t \mathbf{x}_0 + \sigma_t^2 \boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. But in general, $\mathcal{R}(\alpha_t \mathbf{x}_0 + \sigma_t^2 \boldsymbol{\epsilon}) \neq \mathbf{0}$; thus enforcing $\mathcal{R}(\mathbf{x}_t) = \mathbf{0}$ does not align with the underlying marginal distribution.

810 A.2 MEAN ESTIMATION FOR SCORE-BASED MODELS

811
812 As summarized in Appendix A.1, score-based models are an alternative perspective on diffusion
813 models that may be understood as their continuous-time generalization (Song et al., 2021). Here,
814 the model learns the data score $\mathbf{s}_\theta(\mathbf{x}_t, t) \approx \nabla \log p_t(\mathbf{x})$, with which we can solve the reverse SDE
815 equation 15 to generate new samples. For a forward conditional marginal of form $p(\mathbf{x}_t|\mathbf{x}_0) =$
816 $\mathcal{N}(\mathbf{x}_t; \alpha_t \mathbf{x}_0, \sigma_t^2 \mathbf{I})$, we can obtain $\hat{\mathbf{x}}_0 = \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t]$ as (see e.g., Kingma et al. (2021))

$$817 \hat{\mathbf{x}}_0 = \frac{1}{\alpha_t} (\mathbf{x}_t + \sigma_t^2 \mathbf{s}_\theta(\mathbf{x}_t, t)), \quad (18)$$

818 based on which the virtual residual likelihood can be estimated.
819

821 A.3 DETAILS ON THE SIMPLIFIED TRAINING OBJECTIVE

822
823 As shown in Section 3, we arrive at an optimization objective equation 11 that requires sampling over
824 latents from the learned distribution $\mathbf{x}_{1:T} \sim p_\theta(\mathbf{x}_{1:T})$. We simplify this by instead sampling from the
825 available $q(\mathbf{x}_{1:T})$, which can be understood as ignoring the likelihood ratio, as we aim to minimize

$$826 \mathbb{E}_{\mathbf{x}_0 \sim q(\mathbf{x}_0)} [-\log p_\theta(\mathbf{x}_0)] + \mathbb{E}_{\mathbf{x}_{1:T} \sim p_\theta(\mathbf{x}_{1:T})} [-\log q_{\mathcal{R}}(\hat{\mathbf{r}} = \mathbf{0}|\mathbf{x}_0^*(\mathbf{x}_{1:T}))]$$

$$827 \leq \mathbb{E}_{q(\mathbf{x}_{0:T})} \left[\log \frac{q(\mathbf{x}_{1:T}|\mathbf{x}_0)}{p_\theta(\mathbf{x}_{0:T})} \right] + \mathbb{E}_{\mathbf{x}_{1:T} \sim p_\theta(\mathbf{x}_{1:T})} [-\log q_{\mathcal{R}}(\hat{\mathbf{r}} = \mathbf{0}|\mathbf{x}_0^*(\mathbf{x}_{1:T}))]$$

$$828 = \mathbb{E}_{q(\mathbf{x}_{0:T})} \left[\log \frac{q(\mathbf{x}_{1:T}|\mathbf{x}_0)}{p_\theta(\mathbf{x}_{0:T})} \right] + \mathbb{E}_{\mathbf{x}_{1:T} \sim q(\mathbf{x}_{1:T})} \left[-\frac{p_\theta(\mathbf{x}_{1:T})}{q(\mathbf{x}_{1:T})} \log q_{\mathcal{R}}(\hat{\mathbf{r}} = \mathbf{0}|\mathbf{x}_0^*(\mathbf{x}_{1:T})) \right]$$

$$829$$

$$830 \approx \mathbb{E}_{q(\mathbf{x}_{0:T})} \left[\log \frac{q(\mathbf{x}_{1:T}|\mathbf{x}_0)}{p_\theta(\mathbf{x}_{0:T})} \right] + \mathbb{E}_{\mathbf{x}_{1:T} \sim q(\mathbf{x}_{1:T})} [-\log q_{\mathcal{R}}(\hat{\mathbf{r}} = \mathbf{0}|\mathbf{x}_0^*(\mathbf{x}_{1:T}))]$$

$$831$$

$$832 = \mathbb{E}_{q(\mathbf{x}_{0:T})} \left[\log \frac{q(\mathbf{x}_{1:T}|\mathbf{x}_0)}{p_\theta(\mathbf{x}_{0:T})} - \log q_{\mathcal{R}}(\hat{\mathbf{r}} = \mathbf{0}|\mathbf{x}_0^*(\mathbf{x}_{1:T})) \right].$$

$$833$$

$$834$$

$$835$$

$$836$$

837 We justify this simplification by the assumption that optimizing the variational bound will bring
838 $p_\theta(\mathbf{x}_{1:T})$ close to $q(\mathbf{x}_{1:T})$ and thus reduce this bias with ongoing model training. Evaluating the final
839 expression gives us the presented physics-informed diffusion model loss equation 13.

840 A.4 DDIM SAMPLING

841
842 For completeness, we here provide the deterministic DDIM sampling scheme derived by Song et al.
843 (2020) in terms of $\hat{\mathbf{x}}_0$:

$$844 \mathbf{x}_{\tau_{i-1}} = \sqrt{\bar{\alpha}_{\tau_{i-1}}} \hat{\mathbf{x}}_0(\mathbf{x}_{\tau_i}, t) + \sqrt{\frac{1 - \bar{\alpha}_{\tau_{i-1}}}{1 - \bar{\alpha}_{\tau_i}}} \cdot (\mathbf{x}_{\tau_i} - \sqrt{\bar{\alpha}_{\tau_i}} \hat{\mathbf{x}}_0(\mathbf{x}_{\tau_i}, t)). \quad (19)$$

847 Generally, τ is a sub-sequence of the full denoising sequence $[1, \dots, T]$. During training, \mathbf{x}_t is
848 available and we aim to estimate \mathbf{x}_0 in a given number of reduced timesteps, considering a sequence
849 $\tau = [1, \dots, t]$. As we empirically found no notable improvement by providing multiple intermediate
850 timesteps, but this comes with the cost of the additional forward passes, we set $\tau = [1, t]$. Hence, we
851 sample \mathbf{x}_1 from \mathbf{x}_t in one step via equation 19 and then obtain \mathbf{x}_0 by $\hat{\mathbf{x}}_0(\mathbf{x}_1, t = 1)$, resulting in two
852 forward passes of the model. Note that we reduce the impact of worse estimates at noisier timesteps
853 by the increased variance of the virtual residual likelihood, penalizing deviations from $\mathcal{R}(\mathbf{x}_0) = \mathbf{0}$
854 less.

855 A.5 OTHER CONSTRAINT TYPES AND AUXILIARY OPTIMIZATION OBJECTIVES

856 A.5.1 INEQUALITY CONSTRAINTS

857
858 Inequality constraints are also common in physics (e.g., the second law of thermodynamics). Consid-
859 eration of these constraints of type

$$861 h(\mathbf{x}_0) \leq h_{\max} \quad (20)$$

862 can be introduced via $\mathcal{R}_{\text{ineq}} = \text{ReLU}(h(\mathbf{x}_0) - h_{\max}) = \max(0, h(\mathbf{x}_0) - h_{\max})$, which we analogously
863 introduce into the variational loss as the Gaussian $q_{\mathcal{R}_{\text{ineq}}}(\hat{\mathbf{r}}_{\text{ineq}}|\mathbf{x}_0^*(\mathbf{x}_t, t)) = \mathcal{N}(\hat{\mathbf{r}}_{\text{ineq}}; \mathcal{R}_{\text{ineq}}(\mathbf{x}_0^*), \Sigma_t/c)$
with $\hat{\mathbf{r}}_{\text{ineq}} = 0$.

864 A.5.2 OPTIMIZATION OBJECTIVES

865 Optimization objectives of the form

$$866 \min \mathcal{J}(\mathbf{x}_0) \quad (21)$$

867 can also be considered by treating the optimum as a pseudo-observable \hat{j}_{opt} and a similar strategy as
 868 before. In general, however, this optimum is—unlike \hat{r} and \hat{r}_{ineq} —typically unknown. As a remedy,
 869 we may extend the mismatch of the actual and optimal objective by a pseudo-observable $\hat{r}_{\text{opt}} = 0$ and
 870 pose this as a sample from the exponential distribution, as shown by Rixner & Koutsourelakis (2021):
 871

$$872 \mathcal{J}(\mathbf{x}_0) - \hat{j}_{\text{opt}} + \hat{r}_{\text{opt}} \sim \text{Expon}(\lambda), \quad (22)$$

873 where

$$874 \text{Expon}(x; \lambda) = \begin{cases} \lambda e^{-\lambda x} & x \geq 0 \\ 0 & x < 0 \end{cases}. \quad (23)$$

875 By similar reasoning as before, we introduce

$$876 q_{\mathcal{J}}(\hat{r}_{\text{opt}} | \mathbf{x}_0^*(\mathbf{x}_{1:T})) = \prod_{t=1}^T q_{\mathcal{J}}(\hat{r}_{\text{opt}} | \mathbf{x}_0^*(\mathbf{x}_t, t)), \quad (24)$$

$$877 \text{where } q_{\mathcal{J}}(\hat{r}_{\text{opt}} | \mathbf{x}_0^*(\mathbf{x}_t, t)) = \lambda e^{-\lambda(\mathcal{J}(\mathbf{x}_0^*) - \hat{j}_{\text{opt}})},$$

878 and observe that the log-likelihood

$$879 \log q_{\mathcal{J}}(\hat{r}_{\text{opt}} | \mathbf{x}_0^*(\mathbf{x}_{1:T})) = \sum_{t=1}^T \left[\log(\lambda) - \lambda \mathcal{J}(\mathbf{x}_0^*(\mathbf{x}_t, t)) + \lambda \hat{j}_{\text{opt}} \right] \quad (25)$$

880 decouples \hat{j}_{opt} from \mathbf{x}_0^* due to the properties of the exponential distribution. Therefore, knowledge of
 881 \hat{j}_{opt} is not required for training. The loss function is hence extended to

$$882 L_{\text{PIDM-opt}}(\theta) = \mathbb{E}_{t, \mathbf{x}_{0:T}} \left[\lambda_t \|\mathbf{x}_0 - \hat{\mathbf{x}}_0(\mathbf{x}_t, t)\|^2 + \frac{1}{2\Sigma_t} \|\mathcal{R}(\mathbf{x}_0^*(\mathbf{x}_t, t))\|^2 + \lambda \mathcal{J}(\mathbf{x}_0^*(\mathbf{x}_t, t)) \right]. \quad (26)$$

883 Note that we could also couple the parameter λ of the exponential distribution to the denoising
 884 timestep (similar to $\bar{\Sigma}_t$), but choose to keep it constant here. As usual in multi-objective optimization,
 885 trade-offs between the different loss contributions are expected, depending on the relative weighting
 886 (?).
 887

888 A.6 NUMERICAL STUDIES

889 For the presented evaluations, we introduce the mean absolute residual error as follows

$$890 \mathcal{R}_{\text{MAE}}(\mathbf{x}_0) = \frac{1}{M + N} \left(\sum_{i=1}^M |\mathcal{F}_i[\mathbf{x}_0]| + \sum_{j=1}^N |\mathcal{B}_j[\mathbf{x}_0]| \right), \quad (27)$$

891 where i and j denote the i -th PDE and j -th boundary condition constraint, and we consider a total of
 892 M and N such constraints, respectively.
 893

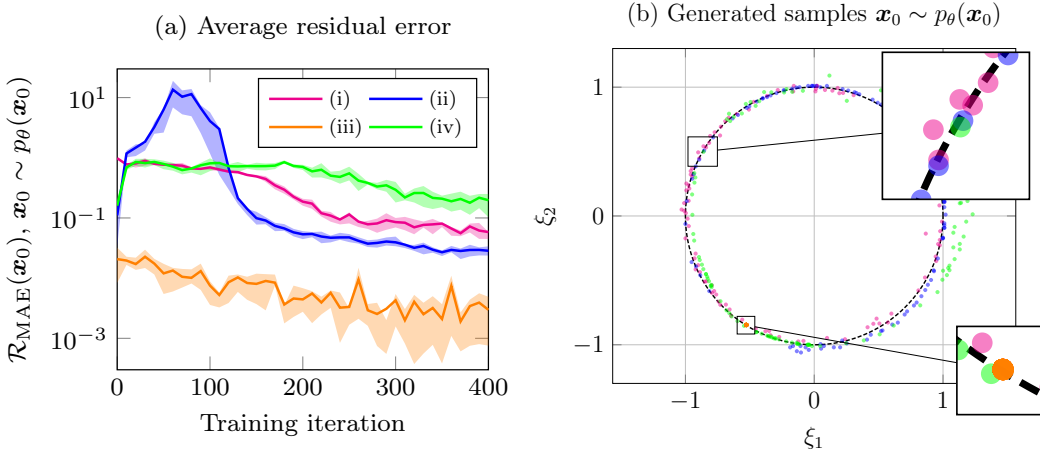
894 A.6.1 TOY PROBLEM

895 **Setup.** In a simple, instructive example, we demonstrate the implications of the proposed physics-
 896 informed loss function equation 13 both considering the mean and sample estimation for the residual
 897 evaluation. The objective is to learn a distribution $q(\mathbf{x}_0)$ that samples points uniformly on the unit
 898 circle $\mathcal{S}_1 = \{\boldsymbol{\xi} \in \mathbb{R}^2 : \|\boldsymbol{\xi}\| = 1\}$, so that samples should obey a simple algebraic equality constraint.
 899 **Thus, generated samples of the model \mathbf{x}_0 directly correspond to the two spatial coordinates (ξ_1, ξ_2) .**
 900 We here identified $c = 0.1$ and $c = 0.005$ via a simple hyperparameter sweep for the mean and
 901 sample estimation, respectively, and thus increased the variance of the implied virtual likelihood.
 902

903 We choose a simple 3-layer MLP of latent dimension 128 with a 2D vector (ξ_1, ξ_2) as in- and output.
 904 Information about the diffusion timestep t is added by transforming t to an embedding which is

918 element-wise multiplied by the output of the linear layer, generally followed by the Softplus activation
 919 except for the last layer. We train the model for 400 epochs on 10,000 randomly sampled points of
 920 the unit circle, using the Adam optimizer (Kingma & Ba, 2014) with a learning rate of 5×10^{-4} . We
 921 use a batch size of 128, and 100 diffusion timesteps with a cosine scheduler (Dhariwal & Nichol,
 922 2021). The training time of each considered model variant takes around 12 minutes on an Nvidia
 923 Quadro RTX 6000 GPU (equipped with 24GB GDDR6 memory).
 924

925 **Results.** To understand the impact of the physics-informed loss, we train the proposed PIDM either
 926 via mean or sample estimation in four different scenarios: (i) via the classical setup of the variational
 927 bound on the data likelihood equation 4, (ii) via the proposed PIDM loss equation 13, (iii) by only
 928 considering the residual loss (and neglecting the variational bound), and (iv) by again considering the
 929 PIDM loss equation 13 but with data sampled from an uninformative prior $\mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. We present
 930 the average residual error of 100 samples generated over 400 training epochs and plot 100 generated
 931 samples after training in Figure 5 and Figure 6 for the mean and sample estimation, respectively.
 932 **More specifically, we provide the mean absolute residual error evaluated as $\mathcal{R}_{\text{MAE}}(\mathbf{x}_0) = \|\|\xi\|^2 - 1\|$,
 933 which is averaged over all samples.**
 934



935
 936
 937
 938
 939
 940
 941
 942
 943
 944
 945
 946
 947
 948
 949
 950 **Figure 5: (Mean estimation)** Evaluation of the average residual error of 100 generated samples
 951 during training (a, averaged over 10 training runs with 25/75%-quantiles using different seeds) and
 952 100 generated samples after training (b, from representative models) in four different settings. We
 953 consider a diffusion model trained with the standard (data-driven) objective (i), our proposed PIDM
 954 (ii), a model trained solely on the residual loss term (iii), and again our proposed PIDM but with data
 955 sampled from an uninformative Gaussian prior (iv). The residual during training is evaluated via
 956 mean estimation, i.e., $\mathbf{x}_0^* = \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t]$. In (b), we indicate the unit circle to which all samples should
 957 be constrained. Colors in (b) match those in (a).

958 We observe that the studies for the mean and sample estimation present a similar picture, and we
 959 may summarize the findings independently of the considered estimation mechanism as follows.
 960 As indicated in Figure 5a and 6a, we observe that the proposed physics-informed loss (ii) indeed
 961 outperforms the standard setup (i) in terms of the residual error after approximately 100-120 training
 962 epochs. This is also visually confirmed in Figure 5b, where the samples from the “constraint-informed”
 963 model align more accurately on the unit circle (see also the top inset). When training the model
 964 solely on the constraints (iii), the model efficiently reduces the residual. However, it converges to
 965 a single point randomly located somewhere on \mathcal{S}_1 , as shown in the bottom inset. This is indeed
 966 expected: the model finds no penalty for collapsing to any points on \mathcal{S}_1 , and similar results are
 967 generally observable if the residual penalty vastly exceeds the data loss term. Interestingly, the
 968 model can also approximate the target distribution even in the absence of training data using only an
 969 uninformative prior (iv), though here at a reduced accuracy. We do not explore this idea further, but
 970 this suggests the use of an uninformative prior as a regularization that prompts the model to explore
 971 a wider range of solutions within the constraint space. Although the simple setting of this study
 prevents straightforward analogies to more complex setups, it confirms the validity of our proposed
 framework and its ability to enforce constraints.

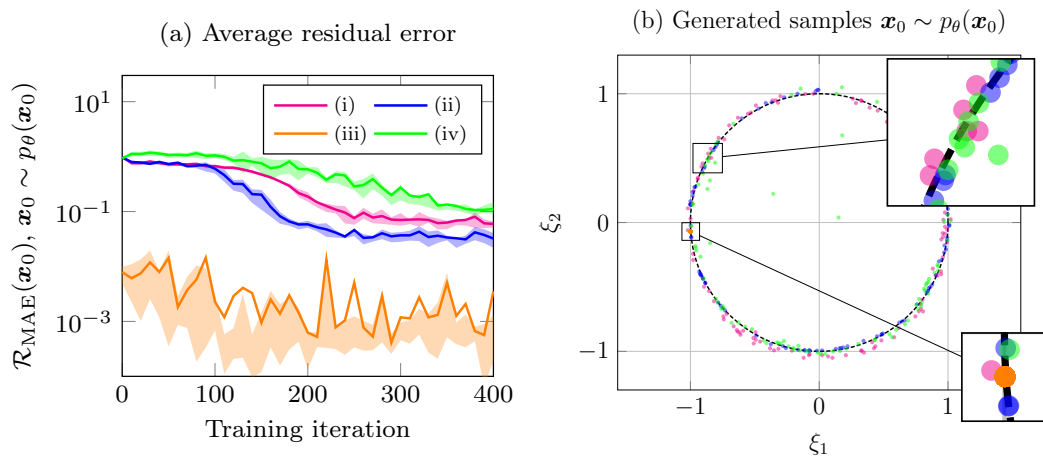


Figure 6: **(Sample estimation)** The setting is identical to the one provided in Figure 5, but we here estimate x_0^* via DDIM (Song et al., 2020).

Study on the influence of relative weighting. To systematically examine the effect of the scale factor c that dictates the importance the model assigns to the residual loss, we provide results for six different values in Figure 7. The experimental setup remains consistent with the previous configuration for the sample estimation, except for the variation in c . As $c \rightarrow 0$, we recover the standard (data-driven) diffusion model. For increasing c , the residual error is increasingly reduced, but the model eventually generates samples increasingly concentrated at a point randomly located on the constraint manifold.

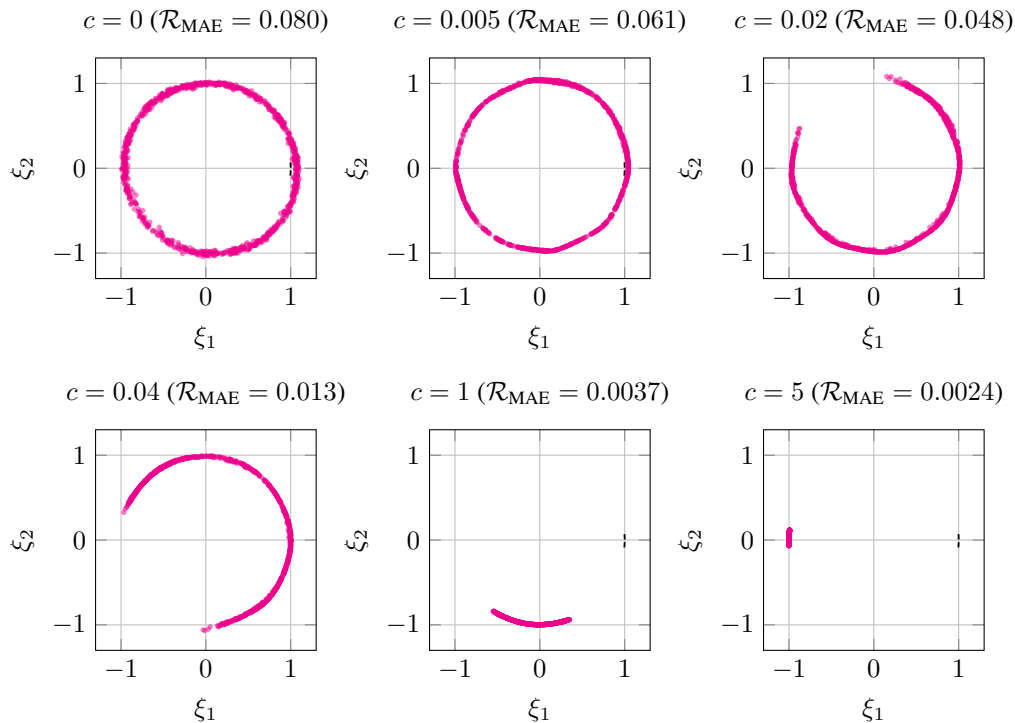


Figure 7: Evaluation of 400 generated samples after training our proposed PIDM with sample estimation in the same setting as Figure 6 but with varying scale factors c . We also indicate the mean absolute residual error \mathcal{R}_{MAE} , averaged over all samples.

As $c \gg 1$, the distribution fully collapses as the model ignores the data loss, and the behavior resembles that of a model trained exclusively on the residual loss (similar to scenario (iii) in Figure 6). For moderate values, such as $c = 0.005$, the distribution maintains full diversity while reducing the residual error, which is the ideal scenario.

A.6.2 DETAILS OF THE DARCY FLOW STUDY

Background. The Darcy flow equations describe the steady-state solution of fluid flow through a porous medium. Given a permeability field $K(\boldsymbol{\xi})$, the pressure distribution $p(\boldsymbol{\xi})$ and velocity field $\mathbf{u}(\boldsymbol{\xi})$ are governed by

$$\begin{aligned} \mathbf{u}(\boldsymbol{\xi}) &= -K(\boldsymbol{\xi})\nabla p(\boldsymbol{\xi}), & \boldsymbol{\xi} \in \Omega \\ \nabla \cdot \mathbf{u}(\boldsymbol{\xi}) &= f(\boldsymbol{\xi}), & \boldsymbol{\xi} \in \Omega \\ \mathbf{u}(\boldsymbol{\xi}) \cdot \hat{\mathbf{n}}(\boldsymbol{\xi}) &= 0, & \boldsymbol{\xi} \in \partial\Omega \\ \int_{\Omega} p(\boldsymbol{\xi}) \, d\boldsymbol{\xi} &= 0, \end{aligned} \tag{28}$$

where $\hat{\mathbf{n}}$ denotes the outward unit vector normal to the boundary. Similar to contemporary work (Zhu & Zabaras, 2018; Jacobsen et al., 2023), we consider a 2D square domain and set the source function to

$$f(\boldsymbol{\xi}) = \begin{cases} r, & \text{if } |\xi_i - \frac{1}{2}w| \leq \frac{1}{2}w, \text{ for } i = 1, 2 \\ -r, & \text{if } |\xi_i - 1 + \frac{1}{2}w| \leq \frac{1}{2}w, \text{ for } i = 1, 2 \\ 0, & \text{otherwise,} \end{cases} \tag{29}$$

with $r = 10$ and $w = 0.125$. We sample $K(\boldsymbol{\xi})$ from a Gaussian random field (GRF), i.e.,

$$K(\boldsymbol{\xi}) = \exp(G(\boldsymbol{\xi})), \quad G(\cdot) \sim \mathcal{N}(0, k(\cdot, \cdot)) \tag{30}$$

with covariance

$$k(\boldsymbol{\xi}, \boldsymbol{\xi}') = \exp(-\|\boldsymbol{\xi} - \boldsymbol{\xi}'\|_2 / l), \quad \text{with } l = 0.1. \tag{31}$$

Instead of directly considering the GRF, we reduce its dimensionality by considering its Karhunen-Loève expansion up to $s = 64$ terms,

$$G(\boldsymbol{\xi}) = \sum_{i=1}^s \sqrt{\lambda_i} z_i \phi_i(\boldsymbol{\xi}), \tag{32}$$

with λ_i and $\phi_i(\boldsymbol{\xi})$ as, respectively, the eigenvalues and eigenfunctions of equation 31 sorted by decreasing λ_i , and $z_i \sim \mathcal{N}(0, 1)$. Discretization of equation 28 is achieved via second-order central finite differences, and we refer to Jacobsen et al. (2023) for details. To incorporate the boundary conditions and the integral constraint, we extend the linear system of type $\mathbf{A}\mathbf{p} = \mathbf{f}$ by additionally imposing the above constraints, so that $\mathbf{A} \in \mathbb{R}^{(n^2+4n+1) \times n^2}$, $\mathbf{p} \in \mathbb{R}^{n^2}$, and $\mathbf{f} \in \mathbb{R}^{n^2+4n+1}$ (where $n = 64$). We solve for the over-determined pressure field using the `scipy.linalg.lstsq` (Virtanen et al., 2020) solver with default settings.

Model architecture and residual computation. The considered U-Net (Ronneberger et al., 2015) architecture is based on Wang (2022), given its demonstrated success in learning the denoising process (Dhariwal & Nichol, 2021). Our configuration uses an image input resolution of 64×64 pixels, aligning with the grid resolution of the linear system under consideration. Importantly, we require the residual evaluation of the predicted images $\mathcal{R}(\mathbf{x}_0^*)$, as required in equation 13, to be consistent with the data, as otherwise the optimal data likelihood is in partial conflict with the (virtual) residual likelihood. This is ensured by using the same finite difference stencils as in the dataset creation, essentially reassembling \mathbf{f} , except that we remove the integral constraint, since it can be trivially fulfilled by shifting the predicted pressure field (Jacobsen et al., 2023). Finite difference stencils are implemented via `torch.nn.Conv2D` (Paszke et al., 2019) with a custom kernel, which we can precompute for stencils up to arbitrary order via `findiff` (Baer, 2018).

We here restrict ourselves to an unconditional model, though extensions to conditional generation (Ho & Salimans, 2022) are straightforward. The U-Net has two in- and output channels and is trained to predict the clean signal based on a noisy input, as stated in Section 2.1. Thus, the model is trained to generate pairs (\mathbf{K}, \mathbf{x}) where \mathbf{K} is sampled similar to equation 30, and \mathbf{p} is the corresponding (unique) pressure field that satisfies the Darcy flow equations equation 28. Further details of the model architecture and training hyperparameters are given in Appendix A.7 and the code in the Supplementary Material.

Comparison with other frameworks. Note that our architecture of the physics-guided diffusion model (ii) is not an exact replication of the one given by Shu et al. (2023) due to the different setting, but the proposed conditioning mechanism is closely mimicked, as detailed in the code provided in the Supplementary Material. Jacobsen et al. (2023) presented a strategy to iteratively “correct” the latent variables \mathbf{x}_t and samples \mathbf{x}_0 during inference by applying gradient descent based on the PDE residuals (iii). We first followed the optimal setting proposed in Jacobsen et al. (2023) and applied gradient-based descent $\nabla_{\mathbf{x}_t} \|\mathcal{R}(\mathbf{x}_t)\|_2^2$ with $\epsilon = 2 \times 10^{-4} / \max \nabla_{\mathbf{x}_t} \mathcal{R}(\mathbf{x}_t)$ for the last N steps of the sampling iterations and M additional iterations. We set $M = 25$ and $N = 50$ according to the best-reported results (equal to starting corrections halfway through the sampling). However, we encountered stability issues for the above value of ϵ , likely due to differences in the sampling scheme and fewer considered timesteps compared to Jacobsen et al. (2023), who adopted a score-based perspective. We therefore conducted a parameter sweep to identify the converging results with the best performance, which we identified at $\epsilon = 1 \times 10^{-6} / \max \nabla_{\mathbf{x}_t} \mathcal{R}(\mathbf{x}_t)$. While additional sweeps might yield some further incremental improvement, we observed that different values of ϵ yielded similar results, as long as the updates converged.

A.6.3 DETAILS OF THE TOPOLOGY OPTIMIZATION STUDY

Background. Topology optimization aims to identify a structure with optimal mechanical properties, typically optimal mechanical stiffness. This can be formalized as the minimization of the mechanical compliance C under a set of equality constraints (given by mechanical equilibrium and boundary conditions) and inequality constraints (typically a volume constraint). Under the assumptions of linear elasticity, the problem reads

$$\begin{aligned} \min_{\rho(\boldsymbol{\xi})} \underbrace{\int_{\Omega} \frac{1}{2} \boldsymbol{\sigma}(\boldsymbol{\xi}) : \boldsymbol{\varepsilon}(\boldsymbol{\xi}) \, d\Omega}_{C}, \quad \boldsymbol{\xi} \in \Omega \\ \text{subject to: } \rho(\boldsymbol{\xi}) \in [0, 1], \\ \int_{\Omega} \rho(\boldsymbol{\xi}) \, d\Omega \leq V_{\max}, \\ \nabla \cdot \boldsymbol{\sigma}(\boldsymbol{\xi}) + \mathbf{f}(\boldsymbol{\xi}) = \mathbf{0}. \end{aligned} \tag{33}$$

The last equation implies quasistatic mechanical equilibrium. Here, $\boldsymbol{\varepsilon}$ and $\boldsymbol{\sigma}$ denote the strain and stress tensor fields, respectively, \mathbf{f} is a distributed body force and V_{\max} a given (maximum) volume constraint. We consider a linear elastic material, which couples $\boldsymbol{\varepsilon}$ and $\boldsymbol{\sigma}$ via Hooke’s law, $\boldsymbol{\sigma} = \mathbf{C} : \boldsymbol{\varepsilon}$, where \mathbf{C} is the fourth-order stiffness tensor, and the colon denotes double tensor contraction. For simplicity, we may assume an isotropic material, so \mathbf{C} is characterized by two material constants (e.g., Young’s modulus E and Poisson’s ratio ν). The solution field is given in terms of the displacements $\mathbf{u}(\boldsymbol{\xi})$, from which the strain tensor follows as $\boldsymbol{\varepsilon} = \frac{1}{2}[\nabla \mathbf{u} + (\nabla \mathbf{u})^T]$. Dirichlet boundary conditions are applied as $\mathbf{u} = \bar{\mathbf{u}}$ on the boundary $\partial\Omega_u \subset \Omega$, and traction boundary conditions $\boldsymbol{\sigma} \cdot \hat{\mathbf{n}} = \mathbf{t}$ on $\partial\Omega_t \subset \Omega$, where $\hat{\mathbf{n}}$ denotes the outward unit vector normal to the boundary $\subset \Omega$.

In practice, equation 33 is usually discretized via finite elements. We here consider a regular finite element mesh, based on a 65×65 grid of nodes with (64×64) four-node quadrilateral elements under plane-stress assumptions (with $E = 1$, $\nu = 0.3$). This turns the mechanical equilibrium equation into a linear system of type $\mathbf{K}\mathbf{U} = \mathbf{F}$, where $\mathbf{K} \in \mathbb{R}^{2n^2 \times 2n^2}$ is the global stiffness matrix and $\mathbf{U}, \mathbf{F} \in \mathbb{R}^{2n^2}$ are the global nodal displacement and the external force vectors, respectively, with $n = 65$. Dirichlet boundary conditions are imposed by appropriately modifying \mathbf{K} and \mathbf{F} . Optimized topologies can subsequently be obtained via the Solid Isotropic Material with Penalization (SIMP) method (we refer to ? for details). SIMP considers continuous densities in equation 33 to allow for gradient-based optimization, defining $\mathbf{C}(\boldsymbol{\xi}) = \rho(\boldsymbol{\xi})^p \mathbf{C}^0$, where $p > 1$ promotes binary entries (corresponding to material placement, $\rho = 1$, or void, $\rho = 0$) and \mathbf{C}^0 denotes the material properties of the given isotropic base material. As SIMP often requires many costly finite element analyses to iteratively refine the solution and may get stuck in local minima, deep learning frameworks including generative adversarial networks (Nie et al., 2021) and diffusion models (Mazé & Ahmed, 2023; Giannone et al., 2023) have been explored as alternatives to mitigate some of these challenges.

Model architecture and residual computation. We consider the same U-Net architecture as in Section 4.1 except for the following differences. Similar to Nie et al. (2021), we also provide the von Mises stress and strain energy fields of the unoptimized domain, as well as the boundary conditions (including the loads) and volume fraction to allow for conditioning in addition to the noisy signal of the solution fields (consisting of the two displacement and density fields) to the model. The model aims to reconstruct the clean signal of the two displacement fields $\mathbf{u}_{1,2}$ and the optimal density ρ . Besides, we increase the latent dimension of the U-Net to 128 to have a similar number of overall parameters compared to previous work (Mazé & Ahmed, 2023; Giannone et al., 2023). Lastly, we apply a sigmoid activation after the density field channel output to ensure $\rho \in [0, 1]$. The predicted $\rho(\xi)$ then enters the governing equations via $\mathbf{C}(\xi) = \rho(\xi)\mathbf{C}^0$.

For the residual and compliance computation, using finite differences as in Section 4.1 to assemble equation 33 introduces inconsistencies with the FEM solution (i.e., training data with non-zero residuals), likely due to the sharp transitions in the density field and point-wise introduction of loads. This is problematic, as the minimization of the residual then does not fully correspond to samples from the data distribution, leading to an optimization conflict. We hence treat the pixels as direct representations of the FE grid and apply the stiffness matrix \mathbf{K} to evaluate the residual. As we consider a U-Net with in- and output pixel dimensions of 64×64 , we apply bilinear interpolation to down- or upscale all nodal quantities to a 65×65 grid where necessary. Note that we can obtain the global stiffness matrix efficiently by precomputing the local stiffness matrix (up to a factor depending on the corresponding density $\rho(\xi)$ at the element level) given, e.g., by `SOLIDSpy` (Gómez & Guarrín-Zapata, 2018) and vectorizing the global assembly. Further details of the model architecture and training hyperparameters are given in Appendix A.7 and the code in the Supplementary Material.

Evaluation. We present metrics consistent with Mazé & Ahmed (2023); Giannone et al. (2023), namely, the compliance error is introduced as $\text{CE} = (\mathbf{C}(\mathbf{x}_0) - \mathbf{C}(\mathbf{x}_{0,\text{SIMP}}))/\mathbf{C}(\mathbf{x}_{0,\text{SIMP}})$ and the volume fraction error as $\text{VFE} = |\bar{\rho}(\mathbf{x}_0) - \rho_{\text{target}}|/\rho_{\text{target}}$, where $\bar{\rho}$ is the (binarized) density averaged over the domain. All models use 100 denoising steps to generate a sample. We note that the models provided by Mazé & Ahmed (2023); Giannone et al. (2023) are trained for only 200k iterations but with a 16 times larger batch size (64). Besides, these frameworks require additional overhead due to the generation of auxiliary data and training of potential surrogate models (Mazé & Ahmed, 2023). For the comparison with PG-Diffusion (Shu et al., 2023) and CoCoGen (Jacobsen et al., 2023) we included both the mechanical equilibrium equations and the volume constraint in the residual. For CoCoGen, computing the full Jacobian $\nabla_{\mathbf{x}_t}\mathcal{R}(\mathbf{x}_t)$ (here, $\mathcal{R}(\mathbf{x}_t)$ is vector-valued), which the authors use to scale the residual exceeds the VRAM of our GPU (even with a batch size of 1 and efficient implementation via `torch.func.jacfwd`). We thus considered a constant scaling factor, conducted an extensive hyperparameter study, and report the best results (in terms of the median residual on the in-distribution test set).

A.7 U-NET ARCHITECTURE AND TRAINING DETAILS

As described in Section 4.1 and 4.2 we consider a U-Net-based architecture (Ronneberger et al., 2015). The main model and training hyperparameters are summarized in Tables 2 and 3, respectively. The model is implemented and trained using PyTorch (Paszke et al., 2019). Further details can be found in the code provided in the Supplementary Material.

For the Darcy flow study, training for 300k iterations took approximately 13 hours for the standard diffusion setup, 16 hours for the PIDM with mean estimation, and 22 hours for the PIDM with sample estimation. For the topology optimization study, training took approximately 48 hours for the standard diffusion setup and 54 hours for the PIDM (with sample estimation). All models were trained on a single Nvidia Quadro RTX 6000 GPU equipped with 24GB GDDR6 memory.

Table 2: Denoising diffusion architecture hyperparameters.

Hyperparameter	Value
In-, output channels (<i>Darcy flow</i>)	2, 2
In-, output channels (<i>Topology optimization</i>)	10, 3
ResNet blocks per down- and upsampling pass	2
ResNet block normalization	Group Normalization (Wu & He, 2018)
ResNet block activation function	SiLU (Elfving et al., 2017)
Attention block normalization	Layer Normalization (Ba et al., 2016)
Feature map resolutions (downsampling pass)	$64 \times 64 \rightarrow 32 \times 32 \rightarrow 16 \times 16 \rightarrow 8 \times 8$
Latent dimensions (in feature maps, <i>Darcy flow</i>)	$32 \rightarrow 64 \rightarrow 128 \rightarrow 256$
Latent dimensions (in feature maps, <i>Top. opt.</i>)	$128 \rightarrow 256 \rightarrow 512 \rightarrow 1024$
Attention (Vaswani et al., 2017; Katharopoulos et al., 2020) head dimension	32
Number of attention heads	8

Table 3: Denoising diffusion process and training hyperparameters.

Hyperparameter	Value
Number of diffusion timesteps	100
β_t -scheduler	Cosine schedule (Dhariwal & Nichol, 2021)
Batch size (<i>Darcy flow</i>), mean estimation	64
Batch size (<i>Darcy flow</i>), sample estimation	16
Batch size (<i>Topology optimization</i> , sample estimation)	4
Iterations (<i>Darcy flow</i>)	300k
Iterations (<i>Topology optimization</i>)	600k
Learning rate	10^{-4}
Optimization algorithm	Adam (Kingma & Ba, 2014)
EMA start (iteration)	1,000
Exponential Moving Average (EMA) decay	0.99
Dropout	none

A.8 ADDITIONAL SAMPLES

A.8.1 DARCY FLOW




Figure 8: Additional samples of permeability and pressure fields as well as the corresponding residual error from the proposed PIDM trained on the Darcy flow dataset. Sample (a) and (b) are sampled from a PIDM with mean estimation, while (c) and (d) are sampled from a PIDM with sample (DDIM) estimation.

A.8.2 TOPOLOGY OPTIMIZATION

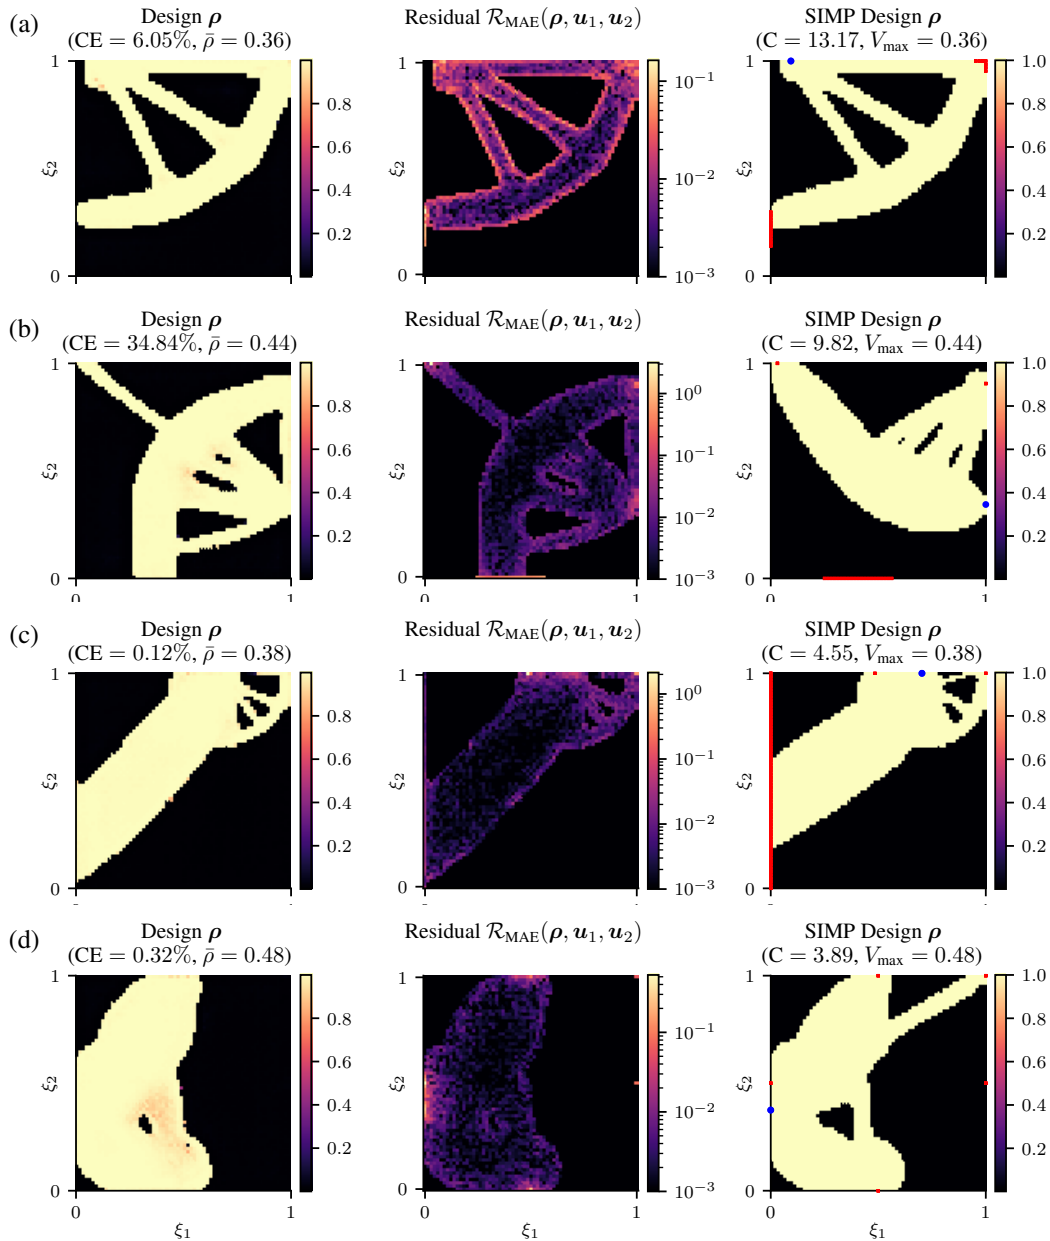


Figure 9: Additional generated designs, including the compliance error CE and volume $\bar{\rho}$, and the residual error (based on the displacement fields, not shown) from the PIDM trained on the SIMP dataset and the corresponding SIMP design, including the compliance C and volume V_{max} . All samples are conditioned on the out-of-distribution test set. In the SIMP design, we indicate the applied load by a blue dot, and the Dirichlet boundary conditions in red.

A.9 EXPLANATION OF THE RESIDUAL

The residual $\mathcal{R}(\mathbf{x}_0)$ serves as a measure of the discrepancy between the generated sample \mathbf{x}_0 and the governing equations it is expected to satisfy. It is defined by the corresponding PDEs and the

1350 respective boundary conditions. More specifically, we stack both contributions in a vector as

$$1351 \quad \mathcal{R}(x_0) := \begin{bmatrix} \mathcal{F}[x_0] \\ \mathcal{B}[x_0] \end{bmatrix}, \quad (34)$$

1354 where

- 1356 • $\mathcal{F}[x_0]$ represents the governing equations, typically a set of PDEs, evaluated over the
- 1357 domain Ω , such that $\mathcal{F}[x_0] = \mathbf{0}$. For instance, in the Darcy flow setting (equation 28), it is
- 1358 assembled as:

$$1359 \quad \mathcal{F}[x_0] := \nabla \cdot (K \nabla p) + f = \mathbf{0}, \quad (35)$$

1360 where K and p correspond to the separate channels of the image x_0 .

- 1361 • $\mathcal{B}[x_0]$ represents the boundary conditions evaluated at the boundary $\partial\Omega$. These conditions
- 1362 might enforce fixed values or specific fluxes at the edges of the domain (see equation 28).

1364 We then obtain a scalar residual loss by considering the squared Euclidean norm of this vector, i.e.,

1365 $\|\mathcal{R}(x_0)\|^2$, as indicated in equation 13.

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403