

GRADIENT-FREE GENERATION FOR HARD-CONSTRAINED SYSTEMS

Anonymous authors

Paper under double-blind review

ABSTRACT

Generative models that satisfy hard constraints are crucial in scientific applications, e.g., numerical simulations, dynamical systems, and supply chain optimization, where physical laws or system requirements must be strictly respected. However, many existing constrained generative models, especially those developed for computer vision, rely heavily on gradient information, which is often sparse or computationally expensive in other fields, e.g., partial differential equations (PDEs). Accurately solving these problems numerically demands the generated solutions to comply with strict physical constraints, e.g., conservation laws. In this work, we introduce a novel framework for adapting pre-trained, unconstrained generative models to exactly satisfy constraints in a zero-shot manner, without requiring expensive gradient computations or fine-tuning. Our framework, ECI sampling, alternates between extrapolation (E), correction (C), and interpolation (I) stages during each iterative sampling step to ensure accurate integration of constraint information while preserving the validity of the generated outputs. We demonstrate the efficacy of our approach across various PDE systems, showing that ECI-guided generation strictly adheres to physical constraints and accurately captures complex distribution shifts induced by these constraints. Empirical results show that our framework consistently outperforms baseline approaches in both zero-shot constrained generative and regression tasks, and achieves competitive results without additional fine-tuning.

1 INTRODUCTION

Diffusion and flow matching models have achieved remarkable success in generative tasks of image generation (Ho et al., 2020; Esser et al., 2024a), language modeling (Lou et al., 2023; Gat et al., 2024), time series prediction (Lin et al., 2024; Kollovieh et al., 2024), and functional data modeling (Lim et al., 2023; Kerrigan et al., 2023). Constrained generation built upon these generative models for solving various inverse problems has also been explored in the image domain (Kawar et al., 2022; Ben-Hamu et al., 2024). These approaches predominantly rely on gradient information with respect to some cost function as a soft constraint. While these soft-constrained methods have been successful in the image domain, applications in other domains often require the generation to adhere to certain constraints exactly. This hard-constrained generation, i.e., generative modeling when the natural constraints are *requirements* and not *suggestions*, is crucial for tasks in many scientific domains. For example, in scientific computing, numerical simulations often require generated solutions to adhere to specific physical constraints (energy or mass conservation (Hansen et al., 2023; Mouli et al., 2024)) and/or satisfy boundary condition constraints (on the values and/or derivatives of the solutions (Saad et al., 2023)).

While existing approaches for constrained generation in image inverse problems can always be adapted (Esser et al., 2024b; Ansari et al., 2024), hard-constrained generation for complex systems like PDE solutions presents the following challenges:

Scarcity of constraint information. Consider common constraints, e.g., the boundary condition (BC) and initial condition (IC) for a partial differential equation (PDE). In contrast to computer vision (CV) tasks, e.g., image inpainting, which typically assume a considerable amount of context information, BC and IC have zero measure with respect to the spatiotemporal domain. For example, a common setting with a spatiotemporal resolution of 100×100 leads to only 1% of known pixel values as context. This is significantly less than the context provided in a typical CV application.

Exact constraint satisfaction. The exact conservation of mass, energy, or momentum is often essential in the simulation of PDEs (LeVeque, 1990). These conservation laws ensure physically feasible and consistent solutions. Similarly, it is often important to exactly adhere to coherency constraints on product hierarchies in supply chain systems. This is contrasted with inverse problems in CV, e.g., super-resolution and de-blurring, where constraints are often implicitly defined as modeling assumptions and evaluation metrics, e.g., peak signal-to-noise ratio (PSNR) also do not directly depend on the exact satisfaction of these constraints. The existing CV approaches fail to guarantee the exact satisfaction of constraints.

Issues with gradient-based methods. Existing zero-shot frameworks for constrained generation predominantly rely on gradient guidance from a differentiable cost function in the final step of generation. This information can be prohibitively expensive to access, e.g., in large 3D PDE spatiotemporal systems; or it may not even exist, e.g., for large systems interacting in complex ways in supply chain systems (Maggiar et al., 2024). Previous work has also indicated drawbacks of these soft-constrained gradient-based approaches including gradient imbalances in the loss terms (Wang et al., 2020; 2021) and ill-conditioning (Krishnapriyan et al., 2021) which can lead to failure modes in scientific machine learning (SciML) tasks.

To address these challenges for hard-constrained generation, we propose a general framework that adopts a gradient-free and zero-shot approach for guiding unconstrained pre-trained flow matching models. For the vast diversity of practical hard constraints in diverse SciML domains, our proposed framework provides a unified and efficient framework without the need for expensive fine-tuning or gradient backpropagation. We term our framework *ECI sampling* since it interleaves *extrapolation*, *correction*, and *interpolation* stages at each iterative sampling step. ECI sampling effectively and accurately captures the distribution shift imposed by the constraints and maintains the consistency of the generative prior. We instantiate ECI sampling on various PDE systems with practical constraints including IC, BC, and conservation laws of physical quantities. We summarize our main contributions as follows:

- **Unified gradient-free generation framework.** We introduce ECI sampling, a unified gradient-free sampling framework for guiding an unconstrained pre-trained flow matching model. By interleaving extrapolation, correction, and interpolation stages at each iterative sampling step, ECI offers fine-grained iterative control over the flow sampling that can accurately capture the distribution shift imposed by the constraints and maintain the consistency of the system.
- **Exact satisfaction of hard constraints.** ECI sampling addresses the unique challenges imposed by hard constraints. The memory- and time-efficient gradient-free approach guarantees the exact satisfaction of constraints and mitigates the gradient issues known with existing gradient-based methods in CV domains.
- **Zero-shot performance on generative and regression tasks.** Comprehensive generative metrics of distributional properties manifest the superior generative performance of our ECI sampling compared with various existing zero-shot guidance methods on various PDE systems. We also show that ECI sampling can be applied to zero-shot regression tasks, still achieving competitive performance with state-of-the-art Neural Operators (NOs) (Li et al., 2020a) that are directly trained on the regression tasks.

2 RELATED WORK

Diffusion and Flow Matching Models. We first review existing generative models for functional data that serve as the generative prior for the spatiotemporal solutions in our approach. Diffusion models (Song & Ermon, 2020; Ho et al., 2020; Song et al., 2020) rely on a variational bound for the log-likelihood and learn a reverse diffusion process to transform prior noise into meaningful samples. Flow matching models (Lipman et al., 2022; Liu et al., 2022) are a family of continuous normalizing flows that learn a time-dependent vector field that defines the data dynamics via the flow ordinary differential equations (ODEs). This flow-based model can be viewed as the continuous generalization of the score matching (diffusion) model that allows for a more flexible design of the denoising process by following the optimal transport formulation (Lipman et al., 2022). Kerrigan et al. (2023) propose functional flow matching (FFM) as an extension of existing flow matching models for image generation. FFM follows the deterministic flow ODE, making it potentially easier to guide than the stochastic Langevin dynamics or stochastic differential equations (SDEs) in the

diffusion formulation. We also note the close connection between diffusion and flow models in Lipman et al. (2022); Albergo et al. (2023). While we focus on guiding flow-based models in this work, we also compare with a wide range of diffusion-based methods.

Constrained Generation. Flow-based generative modeling in function spaces has not been explored until recently, and constrained generation for PDE systems remains largely unexplored. The neural process (NP) (Garnelo et al., 2018; Kim et al., 2019; Sitzmann et al., 2020) is one traditional constrained generation approach, which learns to map a context set of observed input-output pairs to a distribution over regression functions and parameterizes that process in the latent space. Hansen et al. (2023) introduce a generic model-agnostic approach to control the variance of the generation that enjoyed the exact satisfaction of constraints. Lippe et al. (2024) uses gradient guidance from the constraint to guide each sampling step of the pre-trained diffusion model in a zero-shot fashion, similar to diffusion posterior sampling in Chung et al. (2022). Huang et al. (2024) further incorporated the gradients from physics-informed losses (Shu et al., 2023). There are other approaches focused on controlling NOs in regression tasks (Négiar et al., 2023; Saad et al., 2023; Li et al., 2024; Mouli et al., 2024). As the sampling for flow or diffusion models is an iterative procedure, these methods cannot be adopted to provide iterative control. They can be safely incorporated as our *correction* stage method (see Section 3.2).

Inverse Problems. Zero-shot generation for pre-trained diffusion and flow models has been explored in the image domain for solving various inverse problems including image inpainting, deblurring, and superresolution (Bai et al., 2020). Existing approaches predominantly rely on gradient guidance. Liu et al. (2023b); Ben-Hamu et al. (2024); Wang et al. (2024) propose to modify the prior noise by differentiating through the ODE solver. These methods are extremely time- and memory-consuming. Pan et al. (2023a;b) use the adjoint sensitivity method to mitigate the high memory assumption at an even higher sampling time cost. Other works propose gradient-free control, usually with strong prior assumptions on the constraints to derive control at each sampling step. Lugmayr et al. (2022) propose to mix forward diffusion steps from the context with the model’s prediction; and Kawar et al. (2022) use a similar variational approach to solve linear inverse problems.

Although these existing methods have achieved decent generation results for inverse problems in image generation, the approximate enforcement of physical laws (soft-constrained) (Négiar et al., 2023) failed to address the challenges imposed by hard constraints (Wang et al., 2020; 2021; Krishnapriyan et al., 2021). In contrast, our proposed zero-shot guidance approach focuses on the challenges of exact satisfaction of constraints. Our gradient-free approach is also more efficient and provides a unified framework for linear and non-linear constraints with a fine-grain iterative control at each sampling step for a more consistent generation. We summarize the major differences between the existing methods and our proposed method in Table 1.

Table 1: Comparison between existing constrained generation methods and our ECI sampling.

	Zero-shot	Gradient-free	Exact constraint	Iterative control
Conditional FFM [18]	✗	✓	✗	✓
ANP [19]	✓	✓	✗	✗
ProbConserv [13]	✓	✓	✓	✗
DiffusionPDE [16]	✓	✗	✗	✓
D-Flow [4]	✓	✗	✗	✓
ECI (ours)	✓	✓	✓	✓

3 OUR PROPOSED FRAMEWORK: ECI SAMPLING

3.1 PRELIMINARY

Problem Definition. Consider a PDE system $\mathcal{F}_\phi u(x) = 0, x \in \mathcal{X} \subseteq \mathbb{R}^D$ with PDE parameters ϕ . For a parameter set Φ and a fixed PDE family \mathcal{F}_ϕ , let $\mathcal{U}_\mathcal{F} = \{u(x) : \exists \phi \in \Phi, \mathcal{F}_\phi u(x) = 0, x \in \mathcal{X}\}$ denote the set of all the plausible PDE solutions by varying the PDE parameters. Consider some constraint operator $\mathcal{G}u(x) = 0, x \in \mathcal{X}_\mathcal{G} \subseteq \mathcal{X}$ defined on a subset of the PDE domain, and let $\mathcal{U}_\mathcal{G} = \{u(x) : \mathcal{G}u(x) = 0, x \in \mathcal{X}_\mathcal{G}\}$ denote the solution set for the constraint. We are interested in finding the set of solutions $\mathcal{U}_{\mathcal{F}|\mathcal{G}} := \mathcal{U}_\mathcal{F} \cap \mathcal{U}_\mathcal{G} \subseteq \mathcal{U}_\mathcal{F}$ in which both the original PDE operator and the constraint operator are satisfied. We assume that the prior solution set $\mathcal{U}_\mathcal{F}$ can be captured by a generative model, e.g., an FFM pre-trained on a solution set. We aim to guide the pre-trained model to satisfy the constraint operator in a zero-shot fashion towards the narrowed solution set $\mathcal{U}_{\mathcal{F}|\mathcal{G}}$.

Prior Generative Model. FFM (Kerrigan et al., 2023) was proposed to extend the flow matching framework (Lipman et al., 2022) to model measures (distributions) over the Hilbert space of continuous functions. Consider a set \mathcal{U} of “well-behaved” (see Appendix C.2) functions $u : \mathcal{X} \rightarrow \mathbb{R}$, FFM learns a time-dependent vector field operator $v_t : \mathcal{U} \times [0, 1] \rightarrow \mathcal{U}$ that defines a time-dependent diffeomorphism $\psi_t : \mathcal{U} \times [0, 1] \rightarrow \mathcal{U}$ called the *flow* via the differential equation (flow ODE):

$$\partial_t \psi_t(u) = v_t(\psi_t(u)), \quad \psi_0(u) = u_0. \quad (1)$$

The flow ψ_t induces a pushforward measure $\hat{\mu}_t := (\psi_t)_* \mu_0$, where μ_0 is the prior noise measure from which u_0 can be sampled. FFM showed that a tractable flow matching objective can be derived when the flow is conditioned on the target function u_1 sampled from the target measure μ_1 . More details regarding FFM are provided in Appendix C.2 and a more math-

ematically rigorous definition using measure theoretic terms is provided in Appendix A.1 As we are interested in guiding pre-trained generative models, we will assume that FFM can well approximate the target measure $\hat{\mu}_1 \approx \mu_1$ over $\mathcal{U}_{\mathcal{F}}$. We want to guide the pre-trained FFM towards the conditional measure $\mu_{\mathcal{F}|\mathcal{G}}$ over $\mathcal{U}_{\mathcal{F}|\mathcal{G}}$. The sampling procedure of the FFM, similar to all diffusion and flow models, is an iterative process in which the initial noise function is iteratively refined into target functions via the learned vector field with the dynamics described in Equation 1. Algorithm 1 describes sampling from the FFM method using the Euler method. In practice, the vector field can be parameterized by some discretization-invariant NOs (Lu et al., 2019; Li et al., 2020b;a) to make the whole generative framework discretization-invariant. This indicates that FFM can be naturally adapted for zero-shot superresolution (Kerrigan et al., 2023).

Algorithm 1 Sampling from FFM (Euler Method)

- 1: **Input:** Learned vector field v_θ , Euler steps N .
 - 2: Sample noise function $u_0 \sim \mu_0(u)$.
 - 3: **for** $t \leftarrow 0, 1/N, 2/N, \dots, (N-1)/N$ **do**
 - 4: $u_{t+1/N} \leftarrow u_t + v_\theta(u_t, t)/N$
 - 5: **return** u_1
-

3.2 CONSTRAINT GUIDANCE FOR FLOW MODELS

Although various regression NOs have been proposed to satisfy specific constraints in their prediction (Liu et al., 2023a; Négiar et al., 2023; Saad et al., 2023; Liu et al., 2024; Mouli et al., 2024), the iterative nature of flow or diffusion-based models makes it hard to guide the generation toward the exact satisfaction of constraints, as the constraint can only be applied to the last step of generation. Different from existing guidance approaches for diffusion models that predominantly rely on gradient guidance from some loss function of the constraint (soft constraint), we propose a gradient-free method specifically tailored for challenges in the PDE domain.

Following previous zero-shot guidance frameworks (Lugmayr et al., 2022; Kavar et al., 2022), at each iterative sampling timestep t , we model the constrained generation process with the conditional generation probability $p(\hat{u}_t|\mathcal{G}) = p(\hat{u}_t|u_t, \mathcal{G})p(u_t)$, where the unconditional probability $p(u_t)$ can be modeled by a pre-trained generative model. We assume a correction algorithm $C(u_1, \mathcal{G})$ is readily available for *final predictions* orthogonal/oblique projection (Hansen et al., 2023) to ensure hard constraint satisfaction (see Appendix A.3 for concrete correction algorithms). The difficulty in applying these existing algorithms to flow matching models lies in the iterative nature of the flow sampling (Algorithm 1), in which intermediate generations are noised data instead of final predictions, making the correction algorithm inapplicable.

To address this limitation and to offer iterative control over the sampling process, we note that if the final prediction can be *extrapolated* from intermediate noised generations, the correction can be safely applied. Furthermore, if intermediate noised generations can be *interpolated* back from the corrected prediction, we can “propagate” the constraint information back to each time step. Indeed, let $p_\theta(u_1|u_t)$ denote the probability of extrapolating the final prediction of u_1 given current noise data u_t and the learnable parameter θ ; let $q(\hat{u}_t|\hat{u}_1)$ denotes the probability of interpolating the intermediate noised data \hat{u}_t given the corrected data $\hat{u}_1 = C(u_1, \mathcal{G})$, the constraint-conditional probability can be decomposed in a variational way by marginalizing over the auxiliary variable u_1 :

$$p(\hat{u}_t|u_t, \mathcal{G}) = \mathbb{E}_{u_1 \sim p_\theta(u_1|u_t)}[q(\hat{u}_t|C(u_1, \mathcal{G}))] \quad (2)$$

We call such a constraint-guided step in each flow sampling step **ECI Step** (for *extrapolation-correction-interpolation*). Due to deterministic flow ODE formulation along the path of measures $\{\hat{\mu}_t\}$, good properties can be naturally deduced with the flow matching framework. The *extrapolation probability* $p_\theta(u_1|u_t)$ can be learned by the pre-trained unconditional model by doing a deterministic one-step extrapolation of the current predicted vector field as $u_1 = u_t + (1-t)v_\theta(u_t)$. The *interpolation probability* $q(\hat{u}_t|\hat{u}_1)$ is also well-defined in the FFM along the *OT-path* of measures (Kerrigan et al., 2023), where the ground truth data are linearly interpolated with random noises. More specifically, for a sampled noise $u_0 \sim \mu_0(u)$, the interpolation is written as $\hat{u}_t = (1-t)u_0 + t\hat{u}_1$. In this way, the expectation can be discarded and the conditional probability can be further simplified as

$$p(\hat{u}_t|u_t, \mathcal{G}) = q(\hat{u}_t|C(u_t + (1-t)v_\theta(u_t), \mathcal{G})) \quad (3)$$

We summarize the ECI step in Algorithm 2, in which we also allow for $t' \geq t$ to advance the solver. Furthermore, we noted that our ECI steps can be applied recursively to the same timestep, in a manner similar to previous work that applied multiple rounds of guidance (Lugmayr et al., 2022; Ben-Hamu et al., 2024). With the constraint information “backpropagating” into the noise data via one ECI step, a recursive application of such steps can promote further information mixing between the constrained and unconstrained regions, leading to a more consistent generation. At the last mixing step, the interpolation stage instead interpolates the new time step $t' > t$ with sample $u_{t'}$ to advance the ODE solver. The number of total mixing steps M is a controllable hyperparameter. Algorithm 3 describes the complete sampling process for flow matching models using iterative ECI steps, for which we term **ECI sampling**. ECI guarantees exact satisfaction of the constraint and facilitates information mixing between the constrained and unconstrained regions to produce a more consistent generation. See Appendix A for details on the sampling and proof of exact satisfaction.

3.3 CONTROLLING STOCHASTICITY

ECI sampling provides a unified zero-shot and gradient-free guidance framework for various constraints. The varieties in the constraint types may impose challenges. For example, if the constraint contains little information, the variance is expected to be large to match the unconstrained case. On the other hand, with enough constraints to result in a well-posed PDE with a unique solution, the generation variance should be as small as possible (see Section 3.4). For flow-based models, stochasticity only appears in the sampling process of $u_0 \sim \mu_0(u)$.

Inspired by the work of the stochastic interpolant that unified flow matching and diffusion (Albergo et al., 2023), we propose to control the generation stochasticity with different sampling strategies of u_0 during our ECI sampling. Specifically, in the interpolation stage, we sample a random noise function for interpolating the new \hat{u}_t . Previously, we simply use the initially sampled u_0 at $t = 0$. This makes the trajectory straighter as the interpolation is always with respect to the same starting point. An alternative approach is to marginalize over the prior noise measure μ_0 , which results in a more concentrated generation. We use a hyperparameter R to denote the length of the interval for re-sampling the noise function u_0 for interpolation. Intuitively, a smaller R leads to a more concentrated marginalized generation, whereas a larger R leads to larger diversity in the generation.

Algorithm 2 ECISStep

- 1: **Input:** Learned vector field v_θ , constraint \mathcal{G} , current noise data u_t , timestep $t' \geq t$.
 - 2: $u_1 \leftarrow u_t + (1-t)v_\theta$ \triangleright Extrapolation
 - 3: $\hat{u}_1 \leftarrow C(u_1, \mathcal{G})$ \triangleright Correction
 - 4: $u_{t'} \leftarrow (1-t')u_0 + t'\hat{u}_1$ \triangleright Interpolation
 - 5: **return** $u_{t'}$
-

Algorithm 3 ECI Sampling (Euler Method)

- 1: **Input:** Learned vector field v_θ , Euler steps N , mixing iterations M , constraint \mathcal{G} .
 - 2: Sample noise function $u_0 \sim \mu_0(u)$.
 - 3: **for** $t \leftarrow 0, 1/N, 2/N, \dots, (N-1)/N$ **do**
 - 4: $u_t^{(0)} \leftarrow u_t$
 - 5: **for** $m \leftarrow 0, 1, \dots, M-1$ **do**
 - 6: **if** $m < M-1$ **then**
 - 7: $u_t^{(m+1)} \leftarrow \text{ECISStep}(v_\theta, \mathcal{G}, u_t^{(m)}, t)$
 - 8: **else** \triangleright Advance the flow ODE solver
 - 9: $u_{t+1/N} \leftarrow \text{ECISStep}(v_\theta, \mathcal{G}, u_t^{(m)}, t + 1/N)$
 - 10: **return** u_1
-

We quantitatively demonstrate the effect of this hyperparameter in Section 4.3 and also propose heuristic rules for choosing it.

3.4 GENERATIVE FFM AS A REGRESSION MODEL

We illustrate the potential of our ECI sampling framework for various regression tasks. Specifically, the solution set $\mathcal{U}_{|\mathcal{G}}$ defined in Section 3.1 degenerates to a unique solution given enough constraints \mathcal{G} , which leads to a well-posed PDE system. In this way, the generative task degenerates into a deterministic regression learning task, with the target distribution now shifted to the Dirac measure over the solution $\delta_{u_{\text{sol}}}$. Although our constrained generative framework is not designed for these regression tasks and has intrinsic stochasticity, in our empirical study, we show that our framework is able to generate reasonable predictions, compared to standard regression models, e.g., the Fourier Neural Operator (FNO) (Li et al., 2020a).

4 EMPIRICAL EVALUATION

In this section, we present an extensive evaluation of our ECI sampling and other existing zero-shot guidance models for both generative and regression tasks. Covering a wide range of diverse PDE systems, Table 2 summarizes the dataset specifications. See Appendix B for more details regarding the PDE systems and the data generation process. For most zero-shot guidance methods, we first pre-train an unconditional flow matching model as the prior reference model (See Appendix C.3). We test our proposed method using more than one mixing iteration in most cases. See Appendix D for the detailed experimental setup and baseline adaptations for flow sampling, and Appendix E for additional visualizations and results.

Table 2: Dataset specifications (IC / BC / CL stand for initial condition / boundary condition / conservation law). For different resolutions, we test models with the zero-shot superresolution setting.

Type	Dataset	Underspecification	Spatial Resolution	Temporal Resolution	Constraint
Generative	Stokes Problem	IC & BC	100	100	IC / BC
	Heat Equation	IC & diffusion	100	100	IC
	Darcy Flow	BC & field	101×101	N/A	BC
	NS Equation	IC & forcing	64×64	50	IC
Regression	Heat Equation	diffusion	100 / 200	100 / 200	CL
	PME	diffusion	100 / 200	100 / 200	CL
	Stefan Problem	shock range	100 / 200	100 / 200	CL
	Stokes Problem	IC & BC	100	100	IC & BC
	NS Forward	IC & forcing	64×64	50	first 15 frames

4.1 GENERATIVE TASK

In generative tasks, the solution set $\mathcal{U}_{|\mathcal{G}}$ is non-degenerate with an infinite cardinality after imposing the constraint. We expect the generation distribution to match the ground truth shifted distribution. The prior FFM is pre-trained on a larger solution set. During the zero-shot constrained generation, we impose a specific PDE constraint to narrow down the solution set. We use a conditionally-trained FFM (CondFFM), that takes known constraints as conditions during both training and sampling, as an upper bound for the zero-shot guidance models. We conduct extensive experiments on various 2D and 3D PDE datasets with different types of constraints.

To compare two distributions over continuous functions, we follow Kerrigan et al. (2023) to calculate the pointwise mean and standard deviation (std) of the ground truth solutions and the solutions generated from various sampling methods. Mean squared errors (MSE) of the mean (MMSE) and of the standard deviation (SMSE) are used to evaluate the generation quality. To evaluate the satisfaction of the constraint, the constraint error (CE) is defined as the MSE for the constraint operator $\text{CE}(u) = \text{MSE}(\mathcal{G}(u), 0)$ on the constrained region $\mathcal{X}_{\mathcal{G}}$. Inspired by the generative metric of Fréchet inception distance (FID) in CV image generation, we calculate the Fréchet distance between the hidden representations extracted by the pre-trained PDE foundation model Poseidon (Herde et al., 2024), which we term Fréchet Poseidon distance (FPD). A lower FPD indicates a closer match to

Table 3: Generative metrics on various constrained PDEs. The best results for zero-shot methods (CondFFM is not zero-shot) are highlighted in bold.

Dataset	Metric	ECI	ANP	ProbConserv	DiffusionPDE	D-Flow	CondFFM
Stokes IC	MMSE / 10^{-2}	0.090	13.750	2.014	0.601	1.015	0.121
	SMSE / 10^{-2}	0.127	9.183	1.528	0.330	0.219	0.016
	CE / 10^{-2}	0	3.711	0	0.721	1.021	0.017
	FPD	0.076	26.685	13.075	0.202	0.614	0.028
Stokes BC	MMSE / 10^{-2}	0.005	3.003	2.938	0.603	4.800	0.032
	SMSE / 10^{-2}	0.003	3.055	2.302	0.341	217.218	0.047
	CE / 10^{-2}	0	29.873	0	0.007	53.551	0.012
	FPD	0.010	5.339	2.780	1.162	27.680	0.048
Heat Equation	MMSE / 10^{-2}	1.605	5.127	3.744	0.979	1.174	0.008
	SMSE / 10^{-2}	0.157	0.740	4.453	1.291	9.928	0.004
	CE / 10^{-2}	0	0.183	0	11.646	94.551	0.002
	FPD	1.365	2.432	2.639	0.951	3.831	0.006
Darcy Flow	MMSE / 10^{-2}	2.314	14.647	56.608	10.442	1.728	0.043
	SMSE / 10^{-2}	0.592	1.269	65.967	1.097	6.343	0.037
	CE / 10^{-2}	0	3.152	0	0.656	14.553	0.071
	FPD	0.946	5.805	184.910	8.027	10.043	0.021
NS Equation	MMSE / 10^{-2}	7.961	59.020	18.635	19.141	2.846	0.851
	SMSE / 10^{-2}	5.846	14.990	8.306	5.941	6.840	0.554
	CE / 10^{-2}	0	5.527	0	2.295	0.776	0.079
	FPD	1.131	57.263	2.171	5.750	1.651	0.232

the ground truth distribution. To alleviate the impact of randomness in sampling, we sample 512 solutions for the 2D datasets and 100 solutions for the 3D datasets for all methods to calculate the metrics.

Stokes Problem. The 1D Stokes problem from White & Majdalani (2006) is given by the Navier-Stokes equation in the x -coordinate as $u_t = \nu u_{xx}$, with viscosity $\nu \geq 0$. We follow Saad et al. (2023) to construct the dataset varying both IC and BC for pre-training. During constrained sampling, we test two different settings of prescribing fixed IC or BC values, respectively.

Heat Equation. The 1D heat equation is given by the diffusion equation $u_t = \alpha u_{xx}$ with periodic boundary conditions and a sinusoidal IC. We sample the diffusion coefficient α and different ICs for pre-training and fix the IC in the constrained generation.

Darcy Flow. The 2D Darcy flow is a time-independent, second-order, elliptic PDE with the form $-\nabla \cdot (k(x) \nabla u(x)) = f(x)$, $x \in [0, 1]^2$, where k denotes the permeability field and f denotes the forcing function. We follow Li et al. (2020a) to fix $f(x) \equiv 1$ and sample k from a Gaussian random field. We also randomize the constant value imposed as the Dirichlet boundary condition. During constrained sampling, we fix a specific boundary condition.

Navier-Stokes Equation. The 2D Navier-Stokes equation for a viscous, incompressible fluid in 2D can be written for the vorticity w as $\partial_t w(x, t) + u(x, t) \cdot \nabla w(x, t) = \nu \Delta w(x, t) + f(x)$, $x \in [0, 1]^2$, with u determined by w through the Biot-Savart law (Majda et al., 2002). We impose periodic boundary conditions. To compute the ground truth statistics for comparison, we sample 100 initial vorticities and 100 random forces and mesh-grid them to give 10000 training points. 10 different initial vorticities are sampled with the same 100 forces in the training set to form the test set.

Table 3 summarizes the results for the generative tasks with the best zero-shot method performance highlighted in bold. Our proposed ECI sampling achieves state-of-the-art performance on most datasets and also enjoys zero constraint errors. Noticeably, ECI sampling excels in capturing the higher-order distributional properties, e.g., the standard deviation and FPD, and even surpasses the conditional FFM in several simple 2D cases. Figure 1 shows the visualization of point-wise errors for the generation statistics for the Stokes problem with IC fixed together with the ground truth reference. Figure 2 shows the visualization for the NS equation with 25 downsampled time frames. We see that ECI better captures the distributional properties of the shifted conditional mea-

Table 4: Generation setup (batch size vs Euler step), time per sample in seconds, and GPU memory for 2D (Stokes IC) and 3D (NS equation) PDE systems. A fraction sample size is available for ANP. Unfavorable values are highlighted in red.

Dataset	Resource	ECI-1	ECI-5	CondFFM	ANP	ProbConserv	DiffusionPDE	D-Flow
Stokes IC	#sample/#Euler	128/200	128/200	128/200	32/NA	128/200	128/200	2/200
	Time/sample/s	0.065	0.325	0.057	0.009	0.058	0.131	28.774
	GPU Memory/GB	5.4	5.4	5.4	7.4	5.4	10.8	26.4
NS equation	#sample/#Euler	25/100	25/100	25/100	0.2/NA	25/100	25/100	1/20
	Time/sample/s	0.415	2.067	0.669	2.324	0.675	0.676	8.456
	GPU Memory/GB	16.3	16.3	16.3	11.0	16.3	27.0	27.1

sure. Specifically, we observe noticeable artifacts around the IC for the gradient-based methods, i.e., DiffusionPDE and D-Flow, as they do not provide a guarantee of exact satisfaction of the constraint. For ProbConserv (Hansen et al., 2023), although the constraint is satisfied exactly, the naive non-iterative correction approach leaves a sudden change around the boundary. In addition, Table 4 shows that compared with gradient-based methods, our gradient-free approach enjoys sampling efficiency in both 2D and 3D cases. Noticeably, ECI-1 was able to achieve $\times 440$ acceleration and $\times 310$ memory saving compared to gradient-based D-Flow on the Stokes problem. More visualizations for the other PDE systems and the different generation statistics are provided in Appendix E.1.




Figure 1: Generation mean and standard deviation errors for the Stokes problem with IC fixed. Grey for ≈ 0 error, red for positive error, and blue for negative error. Note the noticeable IC errors (left column) for the gradient-based methods.




Figure 2: Generation mean and standard deviation errors of the trajectories for the NS equation with IC fixed. 25 downsampled time frames are plotted for each method.

4.2 REGRESSION TASK

We additionally experiment with regression scenarios (see Section 3.4), where enough constraints reduce the generative task into the standard regression task in neural operator learning. As the ground truth targets are available in this case, we directly calculate the MSE and other related evaluation metrics and compare them with traditional NOs.

Table 5: Uncertainty quantification metrics on constrained PDEs from the Generalized Porous Medium Equation (GPME) family of equations. All baseline results are taken from Hansen et al. (2023). The best results are highlighted in bold.

Dataset	Metric	ECI	ANP	SoftC-ANP	HardC-ANP	ProbConserv
Heat Equation	CE / 10^{-3}	0	4.68	3.47	0	0
	LL	1.90	2.72	2.40	3.08	2.74
	MSE / 10^{-4}	0.81	1.71	2.24	1.37	1.55
PME	CE / 10^{-3}	0	6.67	5.62	0	0
	LL	2.19	3.49	3.11	3.16	3.56
	MSE / 10^{-4}	0.19	0.94	1.11	0.43	0.17
Stefan Problem	CE / 10^{-2}	0	1.30	1.72	0	0
	LL	3.30	3.53	3.57	2.33	3.56
	MSE / 10^{-3}	1.89	5.38	6.81	5.18	1.89

Table 6: Neural operator learning metrics on two constrained PDEs. We abbreviate OOM for out-of-memory. The results for generative models are highlighted in bold.

Dataset	Metric	ECI	ANP	ProbConserv	DiffusionPDE	D-Flow	FNO
Stokes Problem	MMSE / 10^{-3}	0.050	16.470	103.136	5.839	89.514	0.033
	SMSE / 10^{-2}	0.028	0.033	7.509	0.473	136.936	0
NS Equation	MMSE / 10^{-2}	0.069	OOM	46.673	19.015	1.310	0.380
	SMSE / 10^{-2}	0.002	OOM	28.789	13.914	4.068	0

Uncertainty Quantification. In the uncertainty quantification task, random context points from the true solution are sampled together with the conservation laws to pin down the unique solution. Following Hansen et al. (2023), we experiment with our ECI sampling on the Generalized Porous Medium Equation (GPME) family of equations using identical PDE parameter choices. We evaluate the generation results on three instances of the GPME, i.e., the heat equation, porous medium equation (PME), and Stefan problem, where the additional constraints are various conservation laws. We also follow the paper to report the pointwise Gaussian log-likelihood (LL) based on the sample mean and variance. The results in Table 5 show that our ECI sampling is able to achieve comparable results with ProbConserv (Hansen et al., 2023), which is specifically tailored for these uncertainty quantification tasks. We also notice that, though being a generative model with intrinsic stochasticity from the prior noise distribution, our method is able to produce quite confident predictions with little variance. It surpasses the other baselines in terms of MSE but at the cost of worse LL. See Appendix E.2 for visualizations and more discussions.

Neural Operator Learning. We further test our pre-trained model on the Stokes problem and Navier-Stokes equation for the traditional NO learning task. For the Stokes problem, both BC and IC are prescribed to give a unique solution. For the NS equation, we follow Li et al. (2020a) to consider the regression task from the first 15 frames to the other 35 frames. In addition, a standard FNO (Li et al., 2020a) was trained as the regression baseline. Table 6 summarizes the results. The large amount of context pixels in the NS equation makes it infeasible for the ANP model. Our ECI sampling significantly outperforms gradient-based methods and reaches comparable or even better MSEs than the regression NOs directly trained on these regression tasks. It is interesting to observe that even though ECI sampling is a generative model, it is quite confident about its prediction with a small variance. We also provide a comparison of ECI sampling with different numbers of frames fixed as the constraint in Appendix E.3 to demonstrate how the variance gradually reduces.

4.3 ABLATION STUDY

Similar to most existing controlled generation methods for diffusion or flow models, the number of mixing iterations is a sampling hyperparameter for our ECI sampling to control the degree of information mixing. Intuitively, an insufficient number of mixing iterations may fail to mix the in-

formation between the constrained and the unconstrained regions, leading to more artifacts, whereas an excessively large number of mixing iterations may discard the prior PDE structure learned in the pre-trained generative model. We further test with the length of noise re-sampling interval R , as discussed in Section 3.3, to control stochasticity for different constrained generation tasks.

We test the combination of mixing iterations $M \in \{1, 2, 5, 10, 20\}$, with a re-sampling interval $R \in \{1, 2, 5, 10\}$, and no re-sampling settings on the Stokes problem with the IC or BC prescribed. We sample 512 generations for each setting and calculate the MSE for the mean and standard deviation. Figure 3 illustrates the results together with part of the baselines. For the IC task, $R = \text{None}$ performs the best. For the BC task, on the other hand, $R = \text{None}$ performs worse, with $R = 5$ and $R = 2$ achieving lower MSEs. In both cases, the MSEs do not monotonically decrease with an increasing number of mixing iterations. The mixing iteration has an impact on the final performance in a task-specific manner. We hypothesize that this difference in the performance may be caused by the ground truth variations after imposing a specific constraint. For the Stokes problem, the prior ground truth variance is higher for the BC than for the IC. Prescribing BC values provides more information and reduces the uncertainty, making it easier for the model to be guided. In contrast, prescribing IC values provides less information, which results in a large variance in the shifted distribution and makes it harder to guide the model.

Inspired by the aforementioned observations, we propose the following heuristic rules for choosing the appropriate number of mixing iterations M and re-sample interval R : 1) We usually limit M between 1-10, as a larger number of iterations is both computation-intensive and leads to worse performance; and 2) For easier tasks with lower variance, we choose a smaller M and smaller R ; and for harder tasks with higher variance, we choose a larger M and larger R or no re-sampling. Indeed, we use 1 mixing iteration for the relatively easy tasks on the heat equation and Darcy flow and use 10 mixing iterations for the harder NS equation.




Figure 3: MMSE and SMSE of ECI sampling with mixing iterations M in different colors and re-sampling intervals R in different markers in the Stokes problem with IC (left) or BC (right) fixed. “None” means no re-sampling (the initial noise is used).

5 CONCLUSION

In this work, we present ECI sampling as a unified zero-shot and gradient-free framework for systems with hard constraints, which is applicable across diverse domains with functional data and various constraints for flow-based pre-trained models. Instantiated on PDE systems, our framework obtains superior sampling quality and efficient sampling time, compared to existing zero-shot methods; and it also enjoys the additional benefit as a zero-shot regression model with comparable performance to the traditional regression models. We also highlight the potential of ECI sampling on other domains outside of SciML including supply chain and times series, for which we leave as future work. We further note a limitation of our framework: the number of mixing iterations and the re-sampling interval have a considerable task-specific impact on the final performance. While we have provided empirical evaluations of these hyperparameters and suggested heuristic guidelines for selecting them, further extensive evaluations and theoretical analyses would be beneficial.

REFERENCES

- Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023. 3, 5
- Abdul Fatir Ansari, Lorenzo Stella, Caner Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin Shen, Oleksandr Shchur, Syama Syndar Rangapuram, Sebastian Pineda Arango, Shubham Kapoor, Jasper Zschiegner, Danielle C. Maddix, Hao Wang, Michael W. Mahoney, Kari Torkkola, Andrew Gordon Wilson, Michael Bohlke-Schneider, and Yuyang Wang. Chronos: Learning the language of time series. *arXiv preprint arXiv:2403.07815*, 2024. 1
- Yanna Bai, Wei Chen, Jie Chen, and Weisi Guo. Deep learning methods for solving linear inverse problems: Research directions and paradigms. *Signal Processing*, 177:107729, 2020. 3
- Heli Ben-Hamu, Omri Puny, Itai Gat, Brian Karrer, Uriel Singer, and Yaron Lipman. D-flow: Differentiating through flows for controlled generation. *arXiv preprint arXiv:2402.14017*, 2024. 1, 3, 5, 24
- Ricky TQ Chen and Yaron Lipman. Riemannian flow matching on general geometries. *arXiv preprint arXiv:2302.03660*, 2023. 20
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018. 20
- Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*, 2022. 3
- J.R. Dormand and P.J. Prince. A family of embedded runge-kutta formulae. *Journal of Computational and Applied Mathematics*, 6(1):19–26, 1980. ISSN 0377-0427. 24
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024a. 1
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, Kyle Lacey, Alex Goodwin, Yannik Marek, and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis, 2024b. URL <https://arxiv.org/abs/2403.03206>. 1
- Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J Rezende, SM Eslami, and Yee Whye Teh. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018. 3
- Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and Yaron Lipman. Discrete flow matching. *arXiv preprint arXiv:2407.15595*, 2024. 1
- Derek Hansen, Danielle C Maddix, Shima Alizadeh, Gaurav Gupta, and Michael W Mahoney. Learning physical models that can respect conservation laws. In *International Conference on Machine Learning*, pp. 12469–12510. PMLR, 2023. 1, 3, 4, 8, 9, 16, 18, 19, 20, 23, 26
- Maximilian Herde, Bogdan Raonić, Tobias Rohner, Roger Käppeli, Roberto Molinaro, Emmanuel de Bézenac, and Siddhartha Mishra. Poseidon: Efficient foundation models for pdes. *arXiv preprint arXiv:2405.19101*, 2024. 6, 22
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851, 2020. 1, 2, 21
- Jiahe Huang, Guandao Yang, Zichen Wang, and Jeong Joon Park. DiffusionPDE: Generative PDE-solving under partial observation. *arXiv preprint arXiv:2406.17763*, 2024. 3, 24
- Bahjat Kavar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. volume 35, pp. 23593–23606, 2022. 1, 3, 4

- Gavin Kerrigan, Giosue Miglioni, and Padhraic Smyth. Functional flow matching. *arXiv preprint arXiv:2305.17209*, 2023. 1, 2, 3, 4, 5, 6, 15, 21, 23
- Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, Ali Eslami, Dan Rosenbaum, Oriol Vinyals, and Yee Whye Teh. Attentive neural processes. *arXiv preprint arXiv:1901.05761*, 2019. 3, 23
- Marcel Kollovich, Abdul Fatir Ansari, Michael Bohlke-Schneider, Jasper Zschiegner, Hao Wang, and Yuyang Bernie Wang. Predict, refine, synthesize: Self-guiding diffusion models for probabilistic time series forecasting. *Advances in Neural Information Processing Systems*, 36, 2024. 1
- Aditi Krishnapriyan, Amir Gholami, Shandian Zhe, Robert Kirby, and Michael W Mahoney. Characterizing possible failure modes in physics-informed neural networks. In *Advances in Neural Information Processing Systems*, volume 34, pp. 26548–26560, 2021. 2, 3
- Randall J. LeVeque. *Numerical Methods for Conservation Laws*. Lectures in mathematics ETH Zürich. Birkhäuser Verlag, 1990. 2
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020a. 2, 4, 6, 7, 9, 18, 19, 21, 24
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Graph kernel network for partial differential equations. *arXiv preprint arXiv:2003.03485*, 2020b. 4
- Zongyi Li, Hongkai Zheng, Nikola Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Kamyar Azizzadenesheli, and Anima Anandkumar. Physics-informed neural operator for learning partial differential equations. *ACM/JMS Journal of Data Science*, 1(3):1–27, 2024. 3, 24
- Jae Hyun Lim, Nikola B Kovachki, Ricardo Baptista, Christopher Beckham, Kamyar Azizzadenesheli, Jean Kossaifi, Vikram Voleti, Jiaming Song, Karsten Kreis, Jan Kautz, et al. Score-based diffusion models in function space. *arXiv preprint arXiv:2302.07400*, 2023. 1, 21
- Lequan Lin, Zhengkun Li, Ruikun Li, Xuliang Li, and Junbin Gao. Diffusion models for time-series applications: a survey. *Frontiers of Information Technology & Electronic Engineering*, 25(1): 19–41, 2024. 1
- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022. 2, 3, 4, 16, 17, 20, 21
- Phillip Lippe, Bas Veeling, Paris Perdikaris, Richard Turner, and Johannes Brandstetter. Pde-refiner: Achieving accurate long rollouts with neural pde solvers. In *Advances in Neural Information Processing Systems*, volume 37, 2024. 3
- Ning Liu, Yue Yu, Huaqian You, and Neeraj Tatikola. INO: Invariant neural operators for learning complex physical systems with momentum conservation. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206, pp. 6822–6838. PMLR, 2023a. 4
- Ning Liu, Yiming Fan, Xianyi Zeng, Milan Klöwer, Lu Zhang, and Yue Yu. Harnessing the power of neural operators with automatically encoded conservation laws. In *International conference on machine learning*. PMLR, 2024. 4
- Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022. 2
- Xingchao Liu, Lemeng Wu, Shujian Zhang, Chengyue Gong, Wei Ping, and Qiang Liu. Flowgrad: Controlling the output of generative odes with gradients. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 24335–24344, 2023b. 3
- Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion language modeling by estimating the ratios of the data distribution. *arXiv preprint arXiv:2310.16834*, 2023. 1

- Lu Lu, Pengzhan Jin, and George Em Karniadakis. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv preprint arXiv:1910.03193*, 2019. 4
- Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11461–11471, 2022. 3, 4, 5
- A. Maggiar, L. Dicker, and M. W. Mahoney. Consensus planning with primal, dual, and proximal agents. *arXiv preprint arXiv:2408.16462*, 2024. 2
- Andrew J Majda, Andrea L Bertozzi, and A Ogawa. Vorticity and incompressible flow. cambridge texts in applied mathematics. *Appl. Mech. Rev.*, 55(4):B77–B78, 2002. 7
- S. Chandra Mouli, Danielle C. Maddix, Shima Alizadeh, Gaurav Gupta, Andrew Stuart, Michael W. Mahoney, and Yuyang Wang. Using uncertainty quantification to characterize and improve out-of-domain learning for pdes. In *International Conference on Machine Learning*. PMLR, 2024. 1, 3, 4
- Geoffrey Négiar, Michael W. Mahoney, and Aditi S. Krishnapriyan. Learning differentiable solvers for systems with hard constraints. In *International Conference on Learning Representations*, 2023. 3, 4
- Jiachun Pan, Jun Hao Liew, Vincent YF Tan, Jiashi Feng, and Hanshu Yan. Adjointdpm: Adjoint sensitivity method for gradient backpropagation of diffusion probabilistic models. *arXiv preprint arXiv:2307.10711*, 2023a. 3
- Jiachun Pan, Hanshu Yan, Jun Hao Liew, Jiashi Feng, and Vincent YF Tan. Towards accurate guided diffusion sampling through symplectic adjoint method. *arXiv preprint arXiv:2312.12030*, 2023b. 3
- M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. ISSN 0021-9991. 24
- Nadim Saad, Gaurav Gupta, Shima Alizadeh, and Danielle C. Maddix. Guiding continuous operator learning through physics-based boundary constraints. In *International Conference on Learning Representations*, 2023. 1, 3, 4, 7, 18
- Dule Shu, Zijie Li, and Amir Barati Farimani. A physics-informed diffusion model for high-fidelity flow field reconstruction. *Journal of Computational Physics*, 478:111972, 2023. 3
- Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems*, volume 33, pp. 7462–7473, 2020. 3
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems*, volume 32, 2019. 21
- Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. In *Advances in Neural Information Processing Systems*, volume 33, pp. 12438–12448, 2020. 2
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016. 22
- Ashish Vaswani. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017. 21
- Hengkang Wang, Xu Zhang, Taihui Li, Yuxiang Wan, Tiancong Chen, and Ju Sun. Dmplug: A plug-in method for solving inverse problems with diffusion models. *arXiv preprint arXiv:2405.16749*, 2024. 3

Sifan Wang, Yujun Teng, and Paris Perdikaris. Understanding and mitigating gradient pathologies in physics-informed neural networks. *arXiv preprint arXiv:2001.04536*, 2020. 2, 3

Sifan Wang, Hanwen Wang, and Paris Perdikaris. Learning the solution operator of parametric partial differential equations with physics-informed deepnets. *Science advances*, 7(40):eabi8605, 2021. 2, 3

Frank M White and Joseph Majdalani. *Viscous fluid flow*, volume 3. McGraw-Hill New York, 2006. 7

A CONSTRAINED SAMPLING DETAILS

In this section, we further discuss additional details of our proposed ECI sampling framework. Our proposed ECI sampling interleaves extrapolation, correction, and interpolation stages at each sampling step for flow-based generative models to enforce the constraint. We will further elaborate on implementation details as follows.




Figure 4: Intermediate results of the extrapolation, correction, and interpolation stages for the Stokes problem with IC fixed. The trajectories demonstrate how the initial random noise is gradually transformed into the controlled generation with decreasing artifacts around the boundary.

A.1 FORMAL DEFINITION OF CONSTRAINED GENERATION

We further elaborate on our problem setup in more rigorous measure-theoretic terms as a generative task. Consider a well-posed PDE system $\mathcal{F}_\phi u(x) = 0, x \in \mathcal{X}$ with PDE parameters $\phi \in \Phi$, then the solution map $\mathcal{T} : \phi \mapsto u(x)$ is a well-defined mapping. For a fixed distribution p_Φ over the parameter ϕ , we can naturally define the measure over the solution set as the push-forward of the solution map: $\mu_{\mathcal{F}} = \mathcal{T}_* p_\Phi$. Now consider some constraint operator $\mathcal{G}u(x) = 0, x \in \mathcal{X}_{\mathcal{G}} \subseteq \mathcal{X}$ defined on a subset of the PDE domain. Denote the solution set for \mathcal{G} as $\mathcal{U}_{\mathcal{G}} = \{u(x) : \mathcal{G}u(x) = 0, x \in \mathcal{X}_{\mathcal{G}}\}$ and the characteristic function over the solution set as $\chi_{\mathcal{G}}$ such that

$$\chi_{\mathcal{G}}(u) = \begin{cases} 1, & u \in \mathcal{U}_{\mathcal{G}} \\ 0, & u \notin \mathcal{U}_{\mathcal{G}} \end{cases} \quad (4)$$

With the non-degenerating assumption that $\int_{\mathcal{U}} \chi_{\mathcal{G}} d\mu_{\mathcal{F}} > 0$ (this is essentially saying that we should have at least one solution), the conditional measure can be obtained as

$$\mu_{\mathcal{F}|\mathcal{G}}(A) := \mu_{\mathcal{F}}(A|\mathcal{G}) = \int_A \chi_{\mathcal{G}} d\mu_{\mathcal{F}} / \int_{\mathcal{U}} \chi_{\mathcal{G}} d\mu_{\mathcal{F}}, \quad \forall A \subseteq \mathcal{U} \quad (5)$$

where \mathcal{U} is the Hilbert space of all “well-behaving” functions (see [Kerrigan et al. \(2023\)](#) for regularity conditions). We assume $\mu_{\mathcal{F}}$ can be approximated well by generative priors like FFM, and we want to guide it towards the conditional measure $\mu_{\mathcal{F}|\mathcal{G}}$. In practice, the ground truth solution map \mathcal{T} can be obtained by either exact analytical solutions or using numerical PDE solvers. With a pre-defined distribution over PDE parameters (often uniform over some interval, see Table 7 and Appendix B for detailed specification), we can easily sample the PDE parameters and apply the solution map to obtain the sampled function for pre-training.

For constrained sampling, instead of directly calculating the conditional measure $\mu_{\mathcal{F}|\mathcal{G}}$, we can sample the PDE parameters from the pull-back probability $p_{\Phi|\mathcal{G}} = \mathcal{T}^* \mu_{\mathcal{F}|\mathcal{G}}$. In practice, we directly assume such a conditional probability distribution $p_{\Phi|\mathcal{G}}$ over the PDE parameters is known. Indeed, in all the experiments in this paper, we simply choose a subset of PDE parameters and apply the solution map to obtain samples as the ground truth constrained solutions.

A.2 EXTRAPOLATION OF SOLUTIONS

The continuous-time and deterministic formulation of the flow-matching sampling makes it possible to apply different step size to advance the solver for the flow ODE. Specifically, at timestep t , we can make a one-step prediction with a step size of $1 - t$ that directly advances the ODE solver to $t = 1$ as $u_1 = u_t + (1 - t)v$. Theoretically, if the optimal transport paths are learned exactly, the flow should be straight (Lipman et al., 2022) and any arbitrary discretization should lead to the same target. Indeed, in the left plot in Figure 4, the extrapolation gives reasonable prediction even at a timestep close to 0. However, the constraint is not necessarily satisfied in this stage.

A.3 SOLUTION CORRECTION

Our ECI sampling scheme is a unified framework for different constraint operators as long as the corresponding correction algorithm at $t = 1$ is readily available. Specifically, we consider a wide range of constraint operators \mathcal{G} of two following categories: 1) *value constraints* that specify exact values for a subset of the PDE domain: $u(x) = g(x), x \in \mathcal{X}_{\mathcal{G}} \subseteq \mathcal{X}$, and 2) *conservation laws* that specify an exact value for some integration over a subset of the domain: $\int_{\mathcal{X}_{\mathcal{G}} \subseteq \mathcal{X}} u(x) dx = a \in \mathbb{R}$. These two classes of constraints encompass a wide range of practical constraints for PDE systems.

Example 1. For $\mathcal{X} = \Omega \times [0, T]$ where $\Omega \in \mathbb{R}^{D-1}$, the follow examples are value constraints:

- *Initial condition (IC):* $u(x, 0) = g(x), x \in \Omega$.
- *Boundary condition (BC):* $u(x, t) = f(t), x \in \partial\Omega, t \in [0, T]$.

Example 2. For $\mathcal{X} = [0, 1] \times [0, T]$, the following examples are conservation laws:

- *Mass conservation:* $\int_0^1 u(x, t) dx = 0, t \in [0, T]$.
- *Periodic boundary condition (PBC):* $u(0, t) = u(1, t), t \in [0, T]$.

As the constraints encountered in this work are linear constraints, we follow Hansen et al. (2023) to apply orthogonal projection as corrections. The corresponding correction algorithms are outlined in Algorithm 4 and 5.

Algorithm 4 Value Constraint Correction

- 1: **Input:** function u_1 , constraint function g , region $\mathcal{X}_{\mathcal{G}}$.
 - 2: $\hat{u}_1 \leftarrow \mathbb{1}[x \in \mathcal{X}_{\mathcal{G}}] \odot g + \mathbb{1}[x \notin \mathcal{X}_{\mathcal{G}}] \odot u_1$
 - 3: **return** \hat{u}_1
-

Algorithm 5 Conservation Law Correction

- 1: **Input:** function u_1 , conservation value a , region $\mathcal{X}_{\mathcal{G}}$.
 - 2: $g \leftarrow u_1 + (a - \int_{\mathcal{X}_{\mathcal{G}}} u_1(x) dx) / \int_{\mathcal{X}_{\mathcal{G}}} dx$
 - 3: $\hat{u}_1 \leftarrow \mathbb{1}[x \in \mathcal{X}_{\mathcal{G}}] \odot g + \mathbb{1}[x \notin \mathcal{X}_{\mathcal{G}}] \odot u_1$
 - 4: **return** \hat{u}_1
-

It is easy to verify that the corrected function \hat{u}_1 in these correction algorithms indeed satisfies the corresponding constraint exactly. We can also have multiple (even infinitely many) non-overlapping constraints. For example, consider the generalized mass conservation laws $\int_0^1 u(x, t) dx = f(t) \equiv 0, t \in [0, T]$, we have a different conservation law for each PDE timestep t . After discretization, we have a finite set of non-overlapping constraints that can be corrected according to Algorithm 5. The IC correction stage is demonstrated in the middle plot in Figure 4. Note how directly applying the correction algorithm creates noticeable artifacts around the boundary and how such artifacts gradually disappear when advancing the ODE solver.

A.4 INTERPOLATION OF SOLUTIONS

The iterative nature of flow sampling requires interpolation back to arbitrary timestep t during each sampling step. This can be understood as the forward noising process along the conditional probability path and can be efficiently calculated as the linear interpolation $u_t = (1-t)u_1 + tu_0$ according to the conditional flow matching formulation (Lipman et al., 2022). In the right plot in Figure 4, we demonstrate how the initial Gaussian process is transformed into the constrained generation. We also discussed the impact of the re-sampling interval length to control stochasticity in the generation in Section 3.3.

A.5 PROOF OF EXACT SATISFACTION OF CONSTRAINTS

Proposition 1. *Suppose the corrected function $C(u_1, \mathcal{G})$ in Equation 2 satisfies the constraint \mathcal{G} exactly, then for any number of mixing steps $M \geq 1$, the ECI sampling scheme described in Algorithm 3 exactly recovers the constraint in the final generation at $t = 1$.*

Proof. As the advancing step for the ODE solver (the last mixing step) will always perform, it suffices to consider $M = 1$. Consider the last Euler step at $t = 1 - 1/N$, the linear interpolation procedure $q(\hat{u}_{t'}|\hat{u}_1) = 0 \cdot u_0 + 1 \cdot \hat{u}_1 = \hat{u}_1$ will be deterministic as the noise will not contribute to the final interpolant at timestep $t' = 1$. Therefore, the interpolation will deterministically produce $\hat{u}_1 = C(p_\theta(u_1|u_{1-1/N}), \mathcal{G})$ which satisfies the constraint exactly. \square

B DATASET DESCRIPTION AND GENERATION

In this section, we provide a more detailed description of the datasets used in this work and their generation procedure. In addition to the statistics in Table 2, we further provide the dataset specifications in Table 7. For dataset types, *synthetic* indicates that the exact solutions are calculated on the fly based on the randomly sampled PDE parameters for both the training and test datasets. A manual assignment of the dataset sizes is applied with 5k solutions for the training set and 1k for the test set. On the other hand, *simulated* indicates that the solutions are pre-generated using numerical PDE solvers and are different for the training and test datasets.

Table 7: More dataset specifications.

Dataset	PDE Parameter	Split	Spatial Domain	Time Domain	Type
Stokes Problem	$k \sim U[2, 20]$ $\omega \sim U[2, 8]$	5k / 1k	$[0, 1]$	$[0, 1]$	Synthetic
Heat Equation	$\alpha \sim U[1, 5]$ $\phi \sim U[0, \pi]$	5k / 1k	$[0, 2\pi]$	$[0, 1]$	Synthetic
Darcy Flow	k (see below) $C \sim U[-2, 2]$	10k / 1k	$[0, 1]^2$	NA	Simulated
NS Equation	w_0, f (see below)	10k / 1k	$[0, 1]^2$	$[0, 49]$	Simulated
PME	$m \sim U[1, 6]$	5k / 1k	$[0, 1]$	$[0, 1]$	Synthetic
Stefan Problem	$u^* \sim U[0.55, 0.7]$	5k / 1k	$[0, 1]$	$[0, 0.1]$	Synthetic

B.1 STOKES PROBLEM

The 1D Stokes problem is given by the heat equation as:

$$\begin{aligned}
 u_t &= \nu u_{xx}, & x &\in [0, 1], t \in [0, 1], \\
 u(x, 0) &= Ae^{-kx} \cos(kx), & x &\in [0, 1], \\
 u(0, t) &= A \cos(\omega t), & t &\in [0, 1].
 \end{aligned} \tag{6}$$

with viscosity $\nu \geq 0$, oscillation frequency ω , amplitude $A > 0$, and $k = \sqrt{\omega/(2\nu)}$. The analytical solution is given by $u_{\text{exact}}(x, t) = Ae^{-kx} \cos(kx - \omega t)$. Note that k and ω independently and

uniquely define the IC and BC, respectively, and fixing both values reduces the PDE to a well-defined deterministic one with a unique solution. We follow [Saad et al. \(2023\)](#) to construct the dataset by fixing $A = 2$ and sampling $\omega \sim U[2, 8]$, $k \sim U[2, 20]$ for pre-training. During constrained sampling, we test two different settings of prescribing IC with $k = 5$ or BC with $\omega = 6$, respectively. In the regression task, we provide both the IC and BC values. Figure 5 shows the ground truth solutions for the unconstrained and constrained Stokes problem.




Figure 5: Ground truth solutions for the Stokes problem. Unconstrained with $\omega \sim U[2, 8]$, $k \sim U[2, 20]$ (left), IC fixed with $k = 5$ (middle), and BC fixed (right) with $\omega = 6$. In the middle figure, all solutions share the same left-column values. In the right figure, all solutions share the same top-row values.

B.2 HEAT EQUATION

The 1D heat (diffusion) equation with periodic boundary conditions is given as

$$\begin{aligned} u_t &= \alpha u_{xx}, & x &\in [0, 2\pi], t \in [0, 1], \\ u(x, 0) &= \sin(x + \varphi), & x &\in [0, 2\pi], \\ u(0, t) &= u(2\pi, t), & t &\in [0, 1], \end{aligned} \quad (7)$$

where α denotes the diffusion coefficient and φ denotes the phase of the sinusoidal IC. The exact solution is given as $u_{\text{exact}}(x, t) = e^{-\alpha t} \sin(x + \varphi)$. We sample $\alpha \sim U[1, 5]$ and $\varphi \sim U[0, \pi]$ for pre-training. During constrained sampling, we fix the phase $\phi = \pi/4$. Figure 6 shows the ground truth solution family.

In the uncertainty quantification task, we follow [Hansen et al. \(2023\)](#) to fix $\phi = 0$ and vary only the diffusion coefficient in $U[1, 5]$ during pre-training. During constrained generation, we use the same setting to fix $\alpha = 1$, $t = 0.5$ for the calculation of the MSE and LL. The global conservation law for the heat equation in Equation 7 is written as

$$\int_0^{2\pi} u(x, t) dx = 0, \quad t \in [0, 1]. \quad (8)$$

B.3 DARCY FLOW

The 2D Darcy flow is a time-independent, second-order, elliptic PDE with the following form:

$$\begin{aligned} -\nabla \cdot (k(x) \nabla u(x)) &= f(x), & x &\in D = [0, 1]^2, \\ u(x) &= C, & x &\in \partial D, \end{aligned} \quad (9)$$

where k denotes the permeability field and f denotes the forcing function. We follow [Li et al. \(2020a\)](#) to fix $f(x) \equiv 1$ and sample $k \sim \psi_* \mathcal{N}(0, (-\Delta + 9I)^{-2})$ with zero Neumann boundary conditions on the Laplacian. The mapping $\psi : \mathbb{R} \rightarrow \mathbb{R}$ takes the value 12 for positive numbers and 3 for the negative numbers and the pushforward is defined pointwise. Differing from previous work, the boundary condition is sampled from $C \sim U[-2, 2]$. We note that for the same k , if u_0 is the solution for prescribing the boundary condition $u(x) = 0, x \in \partial D$, then $u_0 + C$ is the solution when prescribing the boundary condition $u(x) = C, x \in \partial D$. We then simulate 1000 samples with




Figure 6: Ground truth solutions for the heat equation with $\alpha \sim U[1, 5]$, $\varphi \sim U[0, \pi]$.




Figure 7: Ground truth solutions for the Darcy flow with $C \sim U[-2, 2]$. Solution values are unnormalized on the left and normalized per-solution on the right.

different k and the zero boundary condition and sample C on the fly for both the training and testing datasets. During constrained generation, we fix $C = 1$. Figure 7 shows the ground truth solution family.

B.4 NAVIER-STOKES EQUATION

The 2D Navier-Stokes (NS) equation for a viscous, incompressible fluid in the vorticity form with periodic boundary conditions is given as

$$\begin{aligned} \partial_t w(x, t) + u(x, t) \cdot \nabla w(x, t) &= \nu \Delta w(x, t) + f(x), & x \in [0, 1]^2, t \in [0, T], \\ \nabla \cdot u(x, t) &= 0, & x \in [0, 1]^2, t \in [0, T], \\ w(x, 0) &= w_0(x), & x \in [0, 1]^2, \end{aligned} \quad (10)$$

where u denotes the velocity field, $w = \nabla \times u$ denotes the vorticity, and w_0 denotes the initial vorticity. We follow Li et al. (2020a) to sample $w_0 \sim \mathcal{N}(0, 7^{3/2}(-\Delta + 49I)^{-5/2})$ with periodic boundary conditions. The forcing term is $f(x) = 0.1\sqrt{2}\sin(2\pi(x_1 + x_2) + \phi)$, where $\phi \sim U[0, \pi/2]$ and the viscosity is fixed to be $\nu = 10^{-3}$. We sample 100 initial vorticities and 100 forces with linearly spaced ϕ in the interval $[0, \pi/2]$. We mesh-grid the vorticities and forces to build 10000 solutions for the training set. For the testing set, we sample another 10 initial vorticities and meshgrid with the same 100 forces, resulting in 1000 solutions. We use $T = 49$ and sample 50 snapshots of the numerical solutions. For the regression task that maps the first 15 frames to the other 35 frames, we use the same generative NS equation datasets for training. The regression evaluation is done on the first sample on the test set.

B.5 POROUS MEDIUM EQUATION

The nonlinear Porous Medium Equation (PME) with zero initial and time-varying Dirichlet left boundary conditions is given as

$$\begin{aligned} u_t &= \nabla \cdot (u^m \nabla u), & x \in [0, 1], t \in [0, 1], \\ u(x, 0) &= 0, & x \in [0, 1], \\ u(0, t) &= (mt)^{1/m}, & t \in [0, 1], \\ u(1, t) &= 0, & t \in [0, 1], \end{aligned} \quad (11)$$

where $m \geq 1$. The exact solution is given as $u_{\text{exact}}(x, t) = (m \text{ReLU}(t - x))^{1/m}$. We follow Hansen et al. (2023) to sample $m \sim U[1, 5]$ for pre-training and fix $m = 1, t = 0.5$ for sampling. The conservation law for the PME in Equation 11 is written as

$$\int_0^1 u(x, t) dx = \frac{(mt)^{1+1/m}}{m+1}, \quad t \in [0, 1]. \quad (12)$$

Note that this conservation law implicitly contains information about m .

B.6 STEFAN PROBLEM

The Stefan problem is a challenging and nonlinear case of the Generalized Porous Medium Equation (GPME). With fixed Dirichlet boundary conditions, it is given as:

$$\begin{aligned} u_t &= \nabla \cdot (k(u) \nabla u), & x &\in [0, 1], t \in [0, T], \\ u(x, 0) &= 0, & x &\in [0, 1], \\ u(0, t) &= 1, & t &\in [0, T], \\ u(1, t) &= 0, & t &\in [0, T], \end{aligned} \quad (13)$$

where $k(u)$ denotes the nonlinear step function with respect to a fixed shock value u^* :

$$k(u) = \begin{cases} 1, & u \geq u^* \\ 0, & u < u^*. \end{cases} \quad (14)$$

The exact solution is given as

$$u_{\text{exact}}(x, t) = \mathbb{1}[u \geq u^*] \left(1 - (1 - u^*) \frac{\text{erf}(x/(2\sqrt{t}))}{\text{erf}(\alpha)} \right), \quad (15)$$

where $\text{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z \exp(-t^2) dt$ is the error function and α is uniquely determined by u^* as the solution for the nonlinear equation $(1 - u^*)/\sqrt{\pi} = u^* \text{erf}(\alpha) \alpha \exp(\alpha^2)$. We follow Hansen et al. (2023) to sample the shock value $u^* \sim U[0.55, 0.7]$ and use $T = 0.1$. During sampling, we fix $u^* = 0.6, t = 0.05$. The conservation law for the Stefan problem in Equation 13 is

$$\int_0^1 u(x, t) dx = \frac{2(1 - u^*)}{\text{erf}(\alpha)} \sqrt{\frac{t}{\pi}}, \quad t \in [0, 1]. \quad (16)$$

C FUNCTIONAL FLOW MATCHING AS THE GENERATIVE PRIOR

In this section, we provide additional mathematical backgrounds on flow matching, functional flow matching, and our pre-training settings for FFM as the generative prior.

C.1 FLOW MATCHING

We first provide more details regarding the flow matching framework on common Euclidean data (e.g., pixel values of images) before proceeding to functional flow matching. Flow matching (Lipman et al., 2022) is a generative framework built on continuous normalizing flows (Chen et al., 2018). Flow matching tries to learn the time-dependent *vector field* $v_t : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ that defines a continuous time-dependent diffeomorphism called the *flow* $\psi_t : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ via the following *flow ODE*

$$\frac{\partial}{\partial t} \psi_t(x) = v_t(\psi_t(x)), \quad x_0 \sim p_0(x) \quad (17)$$

where p_0 is the initial noise distribution. The flow induces a probability path with the push-forward $p_t = (\psi_t)_* p_0$ for generative modeling. The vanilla flow matching loss can be written as

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{t \sim U[0, 1], x_1 \sim p_1(x)} [\|v_\theta(x_t, t) - u_t(x_t)\|^2] \quad (18)$$

where $x_t := \psi_t(x)$ is the noise data, p_1 is the target data distribution, and $u_t(x_t)$ is the ground truth data vector field at x_t and timestep t . The vanilla flow matching objective is generally intractable, as we do not know the ground truth vector fields. Lipman et al. (2022) demonstrated a key observation that, when the *conditional* probability path is constructed conditioned on the target data $\psi_t(x|x_1)$, the *conditional* vector field $u_t(x_t|x_1)$ can be calculated analytically while sharing the same gradient with the vanilla flow matching objective. The conditional flow matching objective can be written as

$$\mathcal{L}_{\text{CFM}} = \mathbb{E}_{t \sim U[0, 1], x_1 \sim p_1(x)} [\|v_\theta(x_t, t) - u_t(x_t|x_1)\|^2] \quad (19)$$

Following Chen & Lipman (2023), when further conditioned on the noise $x_0 \sim p_0(x)$, the CFM loss can be reparameterized into a simple form:

$$\mathcal{L}_{\text{CFM}} = \mathbb{E}_{t \sim U[0, 1], x_0 \sim p_0(x), x_1 \sim p_1(x)} [\|v_\theta(x_t, t) - (x_1 - x_0)\|^2], \quad x_t = (1 - t)x_0 + tx_1 \quad (20)$$

The linear interpolation above corresponds to the *optimal-transport probability path* (OT-path) in Lipman et al. (2022), which enjoys additional theoretical benefits of straighter vector fields over other options including the variance-preserving path (Ho et al., 2020) or the variance-exploding path (Song & Ermon, 2019). During sampling, the flow ODE in Equation 17 is solved with the learned vector field to obtain the final generation.

C.2 FUNCTIONAL FLOW MATCHING

Functional flow matching (FFM) (Kerrigan et al., 2023) further extends the conditional flow matching framework to modeling functional data — data that are intrinsically continuous. As mentioned in the original work, the major challenge of extending the CFM framework lies in the fact that probability densities are ill-defined over the infinite-dimensional Hilbert space of continuous functions. FFM proposed to generalize the idea of flow matching to define *path of measures* using measure-theoretic formulations. Specifically, consider a real separable Hilbert space \mathcal{U} of functions $u : \mathcal{X} \rightarrow \mathbb{R}$ equipped with the Borel σ -algebra $\mathcal{B}(\mathcal{U})$. FFM learns a time-dependent vector field operator $v_t : \mathcal{U} \times [0, 1] \rightarrow \mathcal{U}$ that defines a time-dependent diffeomorphism $\psi_t : \mathcal{U} \times [0, 1] \rightarrow \mathcal{U}$ called the *flow* via the differential equation

$$\partial_t \psi_t(u) = v_t(\psi_t(u)), \quad u_0 \sim \mu_0(u) \quad (21)$$

where μ_0 is some fixed noise measure from which random continuous functions can be sampled. Similar to CFM, the flow ψ_t induces a push-forward measure $\hat{\mu}_t := (\psi_t)_* \mu_0$ for generative modeling. FFM demonstrated that the conditional formulation in CFM to deduce a tractable flow matching objective can be also adapted for the path of measures under some regularity conditions. In this way, also relying on the optimal transport path of measures, the FFM objective shares a similar format as the CFM loss in Equation 20

$$\mathcal{L}_{\text{FFM}} = \mathbb{E}_{t \sim U[0,1], u_0 \sim \mu_0(u), u_1 \sim \mu_1(u)} [\|v_\theta(u_t, t) - (u_1 - u_0)\|^2], \quad u_t = (1-t)u_0 + tu_1 \quad (22)$$

where $u_t := \psi_t(u|u_1)$ is the interpolation along the *conditional* path of measures and μ_2 is the target data measure. The norm is the standard L^2 -norm for square-integrable functions. Similarly, sampling for FFM can be thought of as solving the flow ODE with the learned vector field. We demonstrate the Euler method for FFM sampling in Algorithm 1.

We also noted that Lim et al. (2023) proposed the diffusion denoising operator (DDO) to as an extension for diffusion models to function spaces. DDO relies on the non-trivial extension of Gaussian measures on function spaces. Compared to DDO, FFM has a more concise mathematical formulation and better empirical generation results. Therefore, we use FFM as our generative prior.

C.3 PRE-TRAINING FOR FFM

We followed Kerrigan et al. (2023) to use Fourier Neural Operator (FNO) (Li et al., 2020a) as the vector field operator parameterization with additional relative coordinates and sinusoidal time embeddings Vaswani (2017) concatenated to the noised function as inputs. Following FFM, for all 2D data (1D PDE with a time dimensional or Darcy flow), the prior noises are sampled from the 2D Gaussian process with a Matérn kernel with a kernel length of 0.001 and kernel variance of 1. For 3D data (2D NS equation), sampling from the Matérn kernel is prohibitively expensive. Though Kerrigan et al. (2023) has indicated that the white noise does not meet the regularity requirement from mathematical considerations, we empirically found that such white noise prior performed well enough. Therefore, for 3D data, we always sample from the standard white noise.

For 2D data, we used a four-layer FNO with a frequency cutoff of 32×32 , a time embedding channel of 32, a hidden channel of 64, and a projection dimension of 256, which gives a total of 17.9M trainable parameters. All FFM were trained on a single NVIDIA A100 GPU with a batch size of 256, an initial learning rate of 3×10^{-4} , and 20k iterations (approximately 1000 epochs for a 5k training dataset).

For 3D data, we used a two-layer FNO with a frequency cutoff of $16 \times 16 \times 16$, a time embedding channel of 16, a hidden channel of 32, and a projection dimension of 256, which gives a total of 9.46M trainable parameters. This model was trained on 4 NVIDIA A100 GPUs with a batch size of 24 (per GPU) for approximately a total number of 2M iterations (or 5000 epochs) with an initial learning rate of 3×10^{-4} .

It is worth noting that our proposed ECI sampling framework, as a unified zero-shot approach for guiding the pre-trained generative model, does not require conditional training like the conditional FFM baseline. Therefore, we only need to train a separate model for each different PDE system, but not for each different constraint. For example, the two generative tasks and the regression task on the Stokes equation were based on the same pre-trained model on the unconstrained Stokes problem dataset.

Table 8: Mean and standard deviation MSEs between the generation by the pre-trained unconstrained FFM model and the test dataset.

Dataset	Stokes Problem	Heat Equation	Darcy Flow	NS Equation
MMSE / 10^{-2}	0.045	0.130	5.421	0.992
SMSE / 10^{-2}	0.039	0.198	4.121	0.371

Table 9: Generative metrics for the prior unconstrained FFM.

Dataset	MMSE / 10^{-2}	SMSE / 10^{-2}	CE / 10^{-2}	FPD
Stokes IC	2.068	1.572	9.792	13.342
Stokes BC	3.695	3.320	177.297	2.824
Heat Equation	3.954	4.737	49.296	2.799
Darcy Flow	58.879	69.432	7.48×10^4	189
NS Equation	18.749	8.423	11.451	2.185

We provide MMSE and SMSE between the generation by the pre-trained unconstrained FFM model and the test dataset in Table 8 as the evaluation metric for our pre-training stage. It can be seen that most FFMs are able to achieve a decent approximation of the prior unconstrained distribution over the solution set with small MMSE and SMSE. For the 2D Darcy flow, the boundary condition has a greater impact on the final solution such that all pixel values should be shifted by the same value, leading to a larger error in the distributional properties. The generation, on the other hand, does seem reasonable.

We also provide the generative metrics of the pre-trained FFM without any guidance as a sanity check in Table 9. Compared to results in Table 3, it can be demonstrated that there is indeed a shift in the distribution from the generative prior learned by the unconstrained FFM, as many errors are large. Our proposed ECI sampling can successfully capture such a shift in the distribution with significantly lower MSEs and a closer resemblance to the ground truth constrained distribution.

D EXPERIMENTAL SETUP

In this section, we provide further details regarding the experimental setups. We also give a brief introduction to the baselines used in our experiments.

D.1 METRICS

As we have discussed in Section 4, for generative tasks, we calculate distributional properties like the mean squared errors of mean and standard deviation (**MMSE** and **SMSE**) to measure the errors from the ground truth mean and standard deviation, respectively. We also use the constraint error (**CE**) to measure the violation of constraints. The Fréchet Poseidon distance (**FPD**) was inspired by the metric of Fréchet Inception distance (FID) widely used in the image generation domain to evaluate the generative quality using the pre-trained InceptionV3 model (Szegedy et al., 2016). By comparing the Fréchet distance between the hidden activations of the data captured by a pre-trained discriminative model, such a score provides a measurement of the similarity between data distributions. A smaller Fréchet distance to the ground truth data distribution indicates a closer resemblance to the ground truth data and thus a better generation quality. Inspired by FID, we proposed to use the pre-trained PDE foundation model *Poseidon* (Herde et al., 2024) to generate the hidden activations. We used the base version of Poseidon with 157.8M model parameters. Poseidon

takes the initial frame, x -velocity field, y -velocity field, pressure field, and timestep as inputs and outputs predictions of the frame after the timestep. We always zeroed out the x -velocity field, y -velocity field, and pressure field in the inputs. For 2D data, the generated solution was directly fed into Poseidon as the initial frames, and the timestep was always fixed to 0. For 3D data, each time frame of the solution trajectories was fed separately into Poseidon. The Fréchet distances were calculated separately for different time frames and were averaged to obtain the final FPD score. As we want to utilize the pre-trained Poseidon to extract high-level representations of the generated solutions, we extract the last hidden activations of the encoder for Fréchet distance calculation. The hidden activation size for the Poseidon base model is $784 \times 4 \times 4$, which we mean-pool into a 784-dimensional vector for FPD calculation.

For regression tasks, we can directly calculate the **MSE** with respect to the ground truth and the **CE** for constraint violation. Following Hansen et al. (2023), the log-likelihood (**LL**) is calculated pointwise as

$$\text{LL} = \log p(u_1 | \mathcal{N}(\hat{\mu}, \hat{\sigma}^2 I)) = -\frac{u_1 - \hat{\mu}}{2\hat{\sigma}^2} - \log \hat{\sigma} - \frac{1}{2} \log(2\pi), \quad (23)$$

where u_1 is the ground truth data and $\hat{\mu}, \hat{\sigma}$ are the generation mean and standard deviation, respectively.

D.2 BASELINE MODELS

CondFFM. The conditional FFM model (Kerrigan et al., 2023) assumes the family of constraints is known *a priori*. During training, the constraints are fed as conditions to the FFM model as additional information. In this way, CondFFM can only handle value-based constraints but not conservation laws. In practice, to handle values of different constrained regions, we copy the input function and set the pixels in the constrained region to the corresponding values and the pixels in the unconstrained region to zero. This condition is then concatenated to the input channel-wise as additional information. Similarly, this condition is fed into the model during the constrained sampling.

We use an almost identical FNO-based model architecture for CondFFM except for one additional channel for the condition. The training hyperparameters are also the same as the corresponding generative FFM for a fair comparison. As CondFFM has additional information on the constraint in both training and sampling, we naturally expect that it shall exhibit better performance than other zero-shot (training-free) baselines. As a drawback for all models that require training adjustment, different CondFFM models need to be trained separately for different constraints even on the same PDE system (e.g., IC and BC for the Stokes problem), making it less flexible to different constraints.

ANP. The Attentive Neural Process (ANP) (Kim et al., 2019) models the conditional distribution of a function u at a specific set of target points $\{x_i\}_{i \in T}$ given another set of context points $\{x_i\}_{i \in C}$. The ANP utilizes the attention mechanism and variational approach to maximize the variational lower bound for the data likelihood. As it is special case of neural processes, the ANP can take both the target points and context points of arbitrary sizes.

In practice, we use 100 random context points and 1000 random target points during training. The model was trained on the same amount of data as the corresponding FFM models. During sampling, the context points are fixed to be the values in the constraints, and all discretized spatiotemporal coordinates were used as the target points in the generation. This ANP model is also zero-shot and can be applied to different value constraints.

ProbConserv. ProbConserv (Hansen et al., 2023) is a general black-box posterior sampling method specifically designed for the exact satisfaction of various PDE constraints. ProbConserv adopts the gradient-free approach that directly projects the generation to the corresponding solution space using the least squares method. When applied to value constraints, ProbConserv (the version that ensures the exact satisfaction of constraints) directly modifies the corresponding pixels to the constraint values, often leaving noticeable artifacts between the constrained and the unconstrained regions. The ProbConserv model in our experiments is built upon the pre-trained FFM model. As ProbConserv directly operates on the final generation, it does not offer fine-grained control over the intermediate steps of the iterative sampling process of flow-based models.

DiffusionPDE. DiffusionPDE (Huang et al., 2024) uses gradient guidance from the constraint loss and the PINN loss (Raissi et al., 2019; Li et al., 2024) to modify the vector field at each step to satisfy the constraint while maintaining the physical structure of the PDE system. This gradient-based approach is also representative of many models for inverse problems. As our PDE system may have variations in the PDE parameters, the PINN loss is not applicable. Thus, we only use the constraint loss to guide the model at each time step. In practice, we follow the original paper to test several guidance strengths from 10^2 to 10^3 and choose the value with the lowest MSEs.

D-Flow. D-Flow (Ben-Hamu et al., 2024) guides the final generation via optimizing the initial noise u_0 rather than modifying the vector field. Similar to DiffusionPDE, assuming the constraint loss $L(\hat{u}_1)$ is differentiable, D-Flow backpropagates the gradient of the loss through the ODE solver to the initial noise as $\nabla_{u_0} L(\hat{u}_1)$, hoping such an optimization can reduce the final constraint loss. This direct gradient-based approach is very memory-demanding and time-consuming, as the simplest Euler solver requires about 100 steps to achieve decent generation results. We followed the original work to use the LBFGS optimizer with a learning rate of 1 and 20 maximum iterations.

FNO. The Fourier Neural Operator (FNO) (Li et al., 2020a) was used as the baseline in regression tasks of neural operator learning. We used an almost identical architecture as the encoder for FFM, except for the missing time embedding part. Similar to CondFFM, the pixels in the constrained region were set to the constraint values with other pixels set to zero. For a fair comparison, the FNO model tried to predict all the pixel values for the solution domain including the constrained ones already given as the input. As the regression models are easier to overfit, we applied early-stopping techniques to choose the best model checkpoint for evaluation. The FNO was trained for 10k iterations for the Stokes problem and 500k iterations for the NS equation when the valid loss plateaued.

D.3 SAMPLING SETUP

The evaluation of generative models has intrinsic stochasticity, as the sampling procedure has randomness in the prior noise distribution. To provide a more robust evaluation of the generated results, we always generate 512 samples for all 2D datasets and all baselines for computing the MSEs, FPDs, and the sampling time in Table 4. 100 samples were generated for all the NS equation tasks. In this way, we aim to minimize the impact of randomness in the generation and provide a better estimation of the statistics and metrics.

For the baseline models, we use the adaptive Dopri5 ODE solver (Dormand & Prince, 1980) for CondFFM and ProbConserv. For ECI and DiffusionPDE, we used 200 Euler steps for 2D datasets and 100 steps for 3D datasets. For D-Flow, due to the limitation of GPU memory, we use 100 Euler steps for 2D datasets and 20 Euler steps for 3D datasets.

E ADDITIONAL RESULTS AND VISUALIZATIONS

In this section, we provide additional experimental results and ablation studies to further demonstrate the effectiveness of our proposed ECI sampling scheme. Additional visualizations of the generations and evaluation metrics are also provided in this section.

E.1 GENERATIVE TASKS

In addition to the error plots in Figure 1 and 2, we further provide visualizations of the generation statistics in Figure 8 (Stokes problem with IC fixed), 9 (Stokes problem with BC fixed), 11 (heat equation), 13 (Darcy flow), and 15 (2D NS equation).

Errors for generation statistics for different systems are also provided in Figure 10 (Stokes problem with BC fixed), 12 (heat equation), and 14 (Darcy flow). It can be demonstrated more clearly that gradient-based methods like DiffusionPDE and D-Flow had more artifacts around the boundary and did not satisfy the constraint exactly. For ProbConserv, though its gradient-free approach indeed ensured the exact satisfaction of the constraint, the non-iterative control of the prior flow model left noticeable artifacts between the constrained and the unconstrained regions. Similar trends can also be observed in other PDE systems.




Figure 8: Generation mean and standard deviation for the Stokes problem with IC fixed.




Figure 9: Generation mean and standard deviation for the Stokes problem with BC fixed.

We also noticed that D-Flow tended to lead high-variance generations. This is probably because a large number of Euler steps leads to an exceedingly complex dynamic that is difficult to optimize. Also, note that directly taking the gradient with respect to the initial noise may break the initial structure sampled from the Matérn kernel and may lead to initial noises never seen by the model during training. These two reasons may also account for the better performance of the D-Flow for the NS equation, in which we used fewer Euler steps and the standard white noises as the prior noise functions.

We further provide additional generative evaluation on the Stokes problem with various ICs/BCs in Table 10 in addition to the results in Table 3. Recall that during unconstrained pre-training, we varied the PDE parameter as $\omega \sim U[2, 8]$, $k \sim U[2, 20]$, and in Table 3 we fixed $k = 5$ or $\omega = 6$. Combined with the results Table 10, it can be clearly demonstrated that our proposed ECI sampling can achieve consistent performance improvement over the baselines in a zero-shot manner with the flexibility to all different IC settings $k = 5, 10, 15$ and BC settings $\omega = 4, 6, 8$. Furthermore, in Table 11, we provide the evaluation results under the zero-shot superresolution setting, where the generation resolution is 200×200 , four times larger than the training set. It is interesting to see that the conditional FFM is the most sensitive method in such a zero-shot setting, probably due to the hard-coded connection of the masked input as the condition.

We also provide the per-frame FPD scores for the NS equation generation results in Figure 16. We noted that ECI sampling tended to have more accurate generations for frames close to the initial




Figure 10: Generation mean and standard deviation errors for the Stokes problem with BC fixed.




Figure 11: Generation mean and standard deviation for the heat equation with IC fixed.




Figure 12: Generation mean and standard deviation errors for the heat equation with IC fixed.

condition, with the FPD scores increasing with the time frames. In gradient-based models of DiffusionPDE and D-Flow, we also noticed a peak in the FPD scores at an early stage which should have been easier for the model. As expected, CondFFM outperformed all zero-shot guidance models by a large margin in such a complex task for almost all frames.

E.2 REGRESSION TASKS

For the uncertainty quantification tasks, the specific snapshot at the test time as described in Appendix B (following the same settings in (Hansen et al., 2023)) is plotted in Figure 17, with the generation mean and 3 times of the generation standard deviation as the confidence interval. Our results share the similar trend described in (Hansen et al., 2023), in which they also observed more uncertainty with the increasing difficulty from the heat equation, PME, to the Stefan problem. Specifically, the solution in the Stefan problem has a shock in the function value, which leads to higher uncertainty around the shock position. For other linear parts of the solution, our ECI sampling generation is able to achieve quite good predictions with low variance and low uncertainty.

For NO learning tasks, we provide additional visualizations of the generation statistics and the reference FNO predictions in Figure 18 (Stokes problem) and 19 (NS equation). It can be demonstrated that, though not directly trained on such regression tasks, our ECI sampling naturally produces generation results with low variance in such scenarios. In comparison, other models often have a




Figure 13: Generation mean and standard deviation for the Darcy flow with BC fixed. Note that the scaling range is different and large variance values are cropped for a better comparison.




Figure 14: Generation mean and standard deviation errors for the Darcy flow with BC fixed. Note that large error values are cropped for a better comparison.




Figure 15: Generation mean and standard deviation for the 2D NS equation with IC fixed.

variance scale similar to those in the generative tasks, making them less suitable for such zero-shot regression tasks.

E.3 MORE ABLATION STUDIES

To provide a more thorough investigation into the impact of mixing iterations M and re-sampling interval R on the final performance, we further calculate the FPD scores in Figure 20 for the two settings for the Stokes problem in addition to Figure 3. The FPD scores are more consistent than the MMSE and SMSE. For the Stokes IC problem with higher variance, generations with no re-sampling work the best, with the FPD scores peaking at 10 mixing iterations. For the Stokes BC problem with lower variance, the best performance combination of hyperparameters has a diagonal




Figure 16: Per frame FPD score for the NS equation with IC fixed.

Table 10: Generative metrics on the Stokes problem with different constraints. The best results for zero-shot methods (CondFFM is not zero-shot) are highlighted in bold.

Dataset	Metric	ECI	ANP	ProbConserv	DiffusionPDE	D-Flow	CondFFM
Stokes IC $k = 10$	MMSE / 10^{-2}	0.113	4.616	1.792	0.512	0.774	0.307
	SMSE / 10^{-2}	0.124	12.400	0.356	0.301	59.377	0.060
	CE / 10^{-2}	0	13.032	0	0.068	3.004	0.011
	FPD	0.156	6.694	3.619	0.340	11.013	0.454
Stokes IC $k = 15$	MMSE / 10^{-2}	0.190	11.148	0.797	0.216	0.088	0.125
	SMSE / 10^{-2}	0.183	3.354	2.751	0.470	12.063	0.010
	CE / 10^{-2}	0	18.278	0	0.194	13.149	0.010
	FPD	0.105	6.509	6.425	0.832	4.745	0.020
Stokes BC $\omega = 4$	MMSE / 10^{-2}	0.676	9.273	3.034	1.488	1.728	0.031
	SMSE / 10^{-2}	0.378	2.777	2.462	2.278	6.343	0.081
	CE / 10^{-2}	0	115.335	0	39.99	14.553	0.012
	FPD	1.671	5.309	1.686	3.954	10.043	0.071
Stokes BC $\omega = 8$	MMSE / 10^{-2}	0.042	4.115	9.008	5.794	3.801	0.028
	SMSE / 10^{-2}	0.026	2.462	2.710	3.288	22.376	0.010
	CE / 10^{-2}	0	73.916	0	160.397	51.056	0.017
	FPD	0.109	6.663	0.843	9.356	13.202	0.111

Table 11: Generative metrics on the Stokes problem with the zero-shot superresolution setting to 200×200 . The best results are highlighted in bold.

Dataset	Metric	ECI	ANP	ProbConserv	DiffusionPDE	D-Flow	CondFFM
Stokes IC	MMSE / 10^{-2}	0.274	14.120	4.338	5.109	OOM	4.242
	SMSE / 10^{-2}	0.515	8.720	2.933	4.167		4.152
	CE / 10^{-2}	0	3.563	0	5.392		12.508
	FPD	2.647	26.965	15.977	19.654		24.798
Stokes BC	MMSE / 10^{-2}	0.863	2.819	5.223	1.610	OOM	4.669
	SMSE / 10^{-2}	0.418	2.870	1.491	1.768		2.706
	CE / 10^{-2}	0	29.168	0	0.068		21.567
	FPD	1.903	5.282	2.021	3.830		9.504

pattern with more mixing iterations pairing best with a smaller re-sampling interval. Such results are also consistent with our heuristic suggestions on choosing the best M and R for different tasks with different variances.

We further provide the ablation studies on the impact of the quality of the prior generative model. Intuitively, a stronger generative prior that better captures the underlying information of the PDE system can be easier to guide. The results are demonstrated in Table 12, with checkpoints from 5k, 10k, and 20k training iterations. A clear trend of better guided-generation performance can be observed in both IC and BC constraints. In this way, the quality of FFM indeed plays an important role in our ECI sampling framework.

Table 12: Generative metrics on the Stokes problem with different constraints using the different FFM checkpoints with different training iterations.

Dataset	Stokes IC			Stokes BC		
#iter	10k	15k	20k	10k	15k	20k
MMSE / 10^{-2}	1.273	0.601	0.090	0.888	0.011	0.005
SMSE / 10^{-2}	1.693	0.177	0.127	0.868	0.015	0.003
FPD	1.863	1.031	0.076	3.101	0.009	0.010




Figure 17: Uncertainty quantification visualization results for three instances of the Generalized Porous Medium Equation (GPME). The 3σ confidence intervals are plotted.




Figure 18: Regression mean and standard deviation for the Stokes problem with both IC and BC fixed. Note that the scaling range for the standard deviation is different from that in the generative task.

We also provide additional ablation studies on ECI sampling behavior as a regression model. Specifically, we explore the regression losses when more information is provided in the constraint. Intuitively, serving as a deterministic regression model, the ECI generation should exhibit fewer errors and less variance with more information provided. In our experimental setup, we fixed the number of mixing iterations to 5 and the re-sampling interval to 1. To control the information in the constraint, we varied the number of initial frames fed into the model as the value constraint from 5 to 25. As our ECI sampling guarantees the exact satisfaction of the constraint, we instead calculated the mean MSE and standard deviation MSE based on the unconstrained region for a fair comparison between different settings. The losses are plotted in Figure 21, in which a clear and almost linear trend in the log scale can be distinguished for both metrics as the number of constrained frames




Figure 19: Regression mean and standard deviation for the NS equation with the first 15 frames as the constraint. Note that the scaling range for the standard deviation is different from that in the generative task.




Figure 20: FPD score of ECI sampling with different mixing iterations and re-sampling intervals in the Stokes problem with IC (left) or BC (right) fixed.

increases. Similarly, the visualization of the mean error and the logarithm of the standard deviation are provided in Figure 22. The model made more accurate predictions with less uncertainty as the number of constrained frames increased. The regions with large mean errors also matched those with large variance (uncertainty). In this sense, our ECI sampling provides a bonus to transform the pre-trained unconditional FFM model into a decent zero-shot uncertainty quantification model with potentially wider applications.




Figure 21: Regression MMSE and SMSE for the NS equation with different numbers of frames fixed as the constraints. Notice a clear and almost log-linear trend for the decrease in both metrics as the number of constrained frames increases.




Figure 22: Regression error of the mean and logarithm of the standard deviation for the NS equation with different numbers of frames fixed as the constraints. We use the same scale for plotting different settings.