ESS-FLOW: TRAINING-FREE GUIDANCE OF FLOW-BASED MODELS AS INFERENCE IN SOURCE SPACE

Anonymous authors

Paper under double-blind review

ABSTRACT

Guiding pretrained flow-based generative models for conditional generation or to produce samples with desired target properties enables solving diverse tasks without retraining on paired data. We present ESS-Flow, a gradient-free method that leverages the typically Gaussian prior of the source distribution in flow-based models to perform Bayesian inference directly in the source space using Elliptical Slice Sampling. ESS-Flow only requires forward passes through the generative model and observation process, no gradient or Jacobian computations, and is applicable even when gradients are unreliable or unavailable, such as with simulation-based observations or quantization in the generation or observation process. We demonstrate its effectiveness on designing materials with desired target properties and predicting protein structures from sparse inter-residue distance measurements.

1 Introduction

In generative modeling, we are given data samples and aim to construct a sampler that approximates the data distribution. Diffusion models (Ho et al., 2020; Song et al., 2021) and continuous normalizing flows (Lipman et al., 2023; Liu et al., 2023; Albergo et al., 2023) achieve this by transporting samples from a simple source distribution to the data distribution. Instead of just unconditional sampling, we often need to generate samples with specific properties or to match observed data in the context of inverse problems. If sufficient paired data is available, we can train conditional generative models (Dhariwal & Nichol, 2021; Rombach et al., 2022; Miller et al., 2024). However, this requires specialized models for each task and large amounts of paired training data.

Training-free conditional generation methods offer a more flexible alternative by reusing pretrained models. Instead of paired training data, these methods often use a likelihood function to measure how well a sample matches the desired observation. Guidance-based methods modify the transport map towards the target distribution using the conditional likelihood score (Song et al., 2021), though the score often has to be approximated. Alternatively, optimization-based methods minimize the negative log-likelihood with gradient descent, using the generative model as either an explicit regularizer through noising-denoising procedures (Martin et al., 2025; Levy et al., 2024) or an implicit one that constraints the gradient flow to the data manifold (Ben-Hamu et al., 2024).

While optimization-based methods are empirically shown to perform well in image inverse tasks, they only provide point estimates rather than samples, offer little guarantees against local optima, and fail when gradients are unreliable or unavailable. Notably, this occurs when the generative process involves non-differentiable operations like quantization, or when likelihood evaluations require non-differentiable simulations which is common in scientific applications (e.g. Alhossary et al., 2015; Alford et al., 2017). Some of the former limitations can be addressed by formulating controlled generation as Bayesian inference, where the pretrained generative model provides the prior, enabling inference methods that sample from the target distribution. However, many state-of-the-art methods for solving Bayesian inverse problems with generative model priors still requires gradients (Chung et al., 2023; Zhang et al., 2025; Janati et al., 2024; Wu et al., 2023) or are limited to linear Gaussian observations (Kelvinius et al., 2025; Cardoso et al., 2024), limiting their applicability.

Our key insight is that many pretrained diffusion and flow-based models can be recast as continuous normalizing flows with Gaussian source distributions. This enables us to perform Bayesian inference in the source space, instead of the complex data space, and use tailored Markov chain Monte Carlo

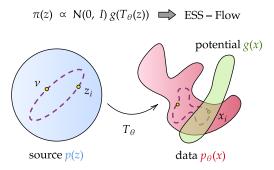


Figure 1: Illustration of ESS-Flow. We sample the target distribution $\pi(z)$ in the source space of flow-based models, which typically has a Gaussian prior p(z). This allows gradient-free sampling with elliptical slice sampling, which defines an ellipse through the current MCMC iterate z_i and a sample from the prior ν , and moves to the next iterate by searching along the ellipse. At stationarity, the transformed ellipse passes through regions of high potential in the data space by construction.

(MCMC) methods like elliptical slice sampling (Murray et al., 2010). This results in ESS-Flow, a new training-free controlled generation method, approximating the target distribution in the source space of flow-based models, as illustrated in Figure 1.

ESS-Flow is by design gradient-free, requiring only point-wise evaluations of the generative model and the likelihood (or more generally, potential) function. This makes it particularly valuable for problems involving quantization, such as molecular and material design with categorical atomic numbers, where gradients are not well-defined. It is applicable for arbitrary, possibly non-linear and non-differentiable, potential functions and does not require knowledge of the noising process used during model training, only the final trained transport map. Furthermore, contrary to guidance-based methods, ESS-Flow preserves the pretrained velocity field. As discussed by Wang et al. (2025), this has the benefit that if the prior vector field is optimized for fast generation, for example when trained with minibatch-OT coupling (Pooladian et al., 2023; Tong et al., 2024), this property is retained by source sampling methods like ESS-Flow.

We summarize our main contributions as:

- We identify and illustrate limitations of gradient-based methods for controlled generation.
- We present ESS-Flow, a training-free and asymptotically exact sampling method for flow-based models, which requires no gradients through the generative model or the potential function.
- We propose a multi-fidelity extension of ESS-Flow, leveraging the fact that flow-based generative models in practice are simulated from using a numerical solver, to improve the computational efficiency of the method.
- We demonstrate ESS-Flow on materials design with target properties and protein structure prediction from partial inter-residue distance measurements, achieving lower mean absolute errors on materials and improved structural realism in proteins.

While the gradient-free nature of ESS-Flow is highly beneficial in many settings, it also limits the applicability of the method in situations when the prior poorly informs the target distribution, for instance when the target is constrained on a lower-dimensional manifold. The primary use-case for ESS-Flow is thus applications, e.g. in scientific domains, where the target distribution is not overly-collapsed. We discuss this limitation further below.

2 Preliminaries

Flow-based generative models: Consider a generative model that transports samples x_0 from a simple source distribution at t=0 to samples x_1 from the data distribution at t=1 through a

learned velocity field $u_t^{\theta}(x)$ by solving the ordinary differential equation (ODE):

$$x_1 = T_{\theta}(x_0) := x_0 + \int_0^1 u_t^{\theta}(x_t) dt.$$
 (1)

Both diffusion models with the probability flow ODE for generation (Song et al., 2021) and models trained with the conditional flow matching objective (Lipman et al., 2023; Albergo et al., 2023) can be viewed as instances of this model class. We collectively refer to them as flow-based generative models with a transport map T_{θ} and consider the case where the source distribution is Gaussian, i.e., the model defines a continuous normalizing flow. To simplify notation, we drop the time subscript t and denote source samples x_0 by t and data samples t by t i.e., the generative model prior is defined by t by t with t collines and t by t collines are the source samples t by t and data samples t by t i.e., the generative model prior is defined by t by t with t collines are the source samples t by t i.e., the generative model prior is defined by t by t in the source samples t by t in the generative model prior is defined by t in the source samples t by t in the generative model prior is defined by t in the source samples t in the source sam

Controlled generation as Bayesian inference: In controlled generation tasks, we seek to sample from a target distribution $\pi(x) \propto g(x)p_{\theta}(x)$, where it is assumed that the normalization constant is finite so that $\pi(x)$ is indeed a probability distribution. Here, $p_{\theta}(x)$ is the prior (unconditional) generative model and g(x) is a nonnegative potential function that guides samples toward desired properties. Under this formulation, $\pi(x)$ is the posterior distribution p(x|y) when g(x) is a likelihood function p(y|x) conditioned on observation y, or a tilted prior when g(x) is a general reward function. In our setting, the prior data distribution is given by the pretrained flow-based generative model of the form in equation (1).

3 RELATED WORK

Many methods have been proposed for controlled generation with diffusion and flow-based models (see surveys in, e.g., Daras et al., 2024; Chung et al., 2025; Zhao et al., 2025). We focus on those that are training-free and applicable to ODE-based generative models of the form in equation (1).

Methods such as DPS (Chung et al., 2023), FlowDPS (Kim et al., 2025), and OT-ODE (Pokle et al., 2024) incorporate a guidance term into the transport map, obtaining target samples in a single pass through the modified generative process. This term involves the score of the conditional likelihood, which is often approximated, and the sequential process cannot correct for errors that occur early in the generation. Methods like PnP-Flow (Martin et al., 2025), DAPS (Zhang et al., 2025), and DDSMC (Kelvinius et al., 2025) overcome this limitation by alternating between noise space and the data space with progressively annealed noise. PnP-Flow optimizes in the data space and regularizes through noising-denoising steps between gradient updates. DAPS generalizes this framework for posterior sampling and DDSMC extends DAPS using sequential Monte Carlo for linear Gaussian potentials. These methods do not require differentiating the transport map and perform approximate posterior updates in the data space at each iteration of the generative procedure. However, they require access to the noising process used during training, as the noising step is part to the algorithm.

Source space methods like D-Flow (Ben-Hamu et al., 2024) performs gradient-based source point optimization, relying on implicit regularization that leads to manifold-constrained gradient flow in the data space. Purohit et al. (2025) extend D-Flow to posterior sampling in the source using Langevin Monte Carlo. In work which is concurrent to ours, Wang et al. (2025) use Hamiltonian Monte Carlo in the source space, which has also been considered earlier by Graham & Storkey (2017) who demonstrated latent space sampling with constrained Hamiltonian Monte Carlo for variational autoencoders. A key limitation of existing source space methods is their reliance on gradients, requiring expensive backpropagation through the ODE solver. To the best of our knowledge, our method, ESS-Flow is the only gradient-free method in this category. Furthermore, it does not require knowledge of the noising process used in training, only the trained transport map.

4 Method

Suppose that we have a generative model prior $p_{\theta}(x)$ from a flow-based model that maps samples from a Gaussian source distribution $z \sim \mathcal{N}(0, I)$ to the data space with the transport map $x = T_{\theta}(z)$. Given a potential function g(x), such as the likelihood p(y|x) for an observation y or a general reward function, we seek samples from the target $\pi(x)$:

$$\pi(x) \propto g(x) p_{\theta}(x) = g(x) p(z) |\det(J_{T_{\theta}}(z))|^{-1},$$
 (2)

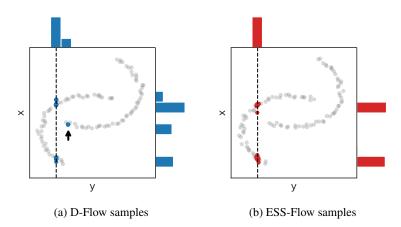


Figure 2: Conditional generation targeting a specific *y* indicated by the dotted line. Prior samples are shown in gray with marginals along the border. Some D-Flow samples, which follow the source-space gradient, become trapped in disconnected manifolds (indicated by arrow), while ESS-Flow, being gradient-free, can avoid this problem.

where $z = T_{\theta}^{-1}(x)$ and $J_{T_{\theta}}(z)$ denotes the Jacobian of the transport map evaluated at z.

4.1 ESS-FLOW: AVOIDING GRADIENT COMPUTATIONS FOR SOURCE SPACE SAMPLING

Sampling the target distribution $\pi(x)$ or optimizing for a point estimate in the data space requires evaluating the Jacobian of the transport map, which is computationally expensive for continuous flow-based models. Methods like PnP-Flow (Martin et al., 2025) avoid this by maximizing only g(x) and regularizing with an iterative noising-denoising procedure, which does not guarantee convergence to the maximum a posteriori (MAP), nor sampling from the target $\pi(x)$.

With a change of variables, we can instead approximate the target distribution $\pi(z)$ in the source:

$$\pi(z) = \pi(x) \left| \det(J_{T_{\theta}}(z)) \right| \propto g(x) \ p(z) \left| \det(J_{T_{\theta}}(z)) \right|^{-1} \left| \det(J_{T_{\theta}}(z)) \right| = g(T_{\theta}(z)) \ p(z). \tag{3}$$

Note the cancellation of the Jacobian terms due to the fact that we express both the prior and the posterior in the source space. While we no longer need the Jacobian for point-wise density evaluation, optimizing for a MAP estimate or using gradient-based sampling with Langevin Monte Carlo (Purohit et al., 2025) or Hamiltonian Monte Carlo (Graham & Storkey, 2017; Wang et al., 2025) still requires it, since the gradient $\nabla_z g(T_\theta(z)) = J_{T_\theta}(z)^\mathsf{T} \nabla_x g(x)$ involves the Jacobian. Beyond computational expense, gradients may be entirely unavailable when the generative process involves quantization as in Section 5.1, or when evaluating the potential requires non-differentiable simulators or external programs that are not amendable to automatic differentiation.

Furthermore, gradient-based methods can struggle when the prior has multiple disconnected modes. To illustrate this, consider D-Flow applied to the toy-problem in Figure 2, where $p_{\theta}(x)$ consists of two interleaving half-circles and we seek samples with a specific target y. D-Flow maximizes $g(T_{\theta}(z))$ w.r.t. z using gradient-based optimization while relying on implicit regularization. A small gradient step $\delta_z \propto J_{T_{\theta}}(z)^{\mathsf{T}} \nabla_x g(x)$ in the source space corresponds to the step $\delta_x = J_{T_{\theta}}(z) J_{T_{\theta}}(z)^{\mathsf{T}} \nabla_x g(x)$ in the data space. This projects gradients along the data manifold, resulting in a manifold-constrained gradient flow (Ben-Hamu et al., 2024). However, as seen in Figure 2, even with a large learning rate that deviates significantly from a gradient flow, many D-Flow samples remain trapped within the disconnected manifold component where they are initialized.

To overcome the limitations of gradient-based methods we propose to use elliptical slice sampling (ESS, Murray et al., 2010) to sample from the target distribution in the source space. This leverages the fact that the prior p(z) is a multivariate Gaussian, whereas the complexity of the transport map $T_{\theta}(z)$ has been incorporated in the potential according to equation (3). This is precisely the setup where ESS excels. Our resulting approach, ESS-Flow, requires only point-wise evaluations of the generative model and involves no gradient or Jacobian computation. Being an MCMC method,

217

218

219

220

221

222

223

224

225

226

227

228

229

230

231

232

233

234235

236

237

238239

240

241

242243244

245

246 247

248

249

250

251

252

253

254

255

256

257258259260

261

262

263

264

265

266267

268

269

ESS-Flow is flexible with initialization, e.g., by sampling z_0 from the Gaussian prior, and then proceeding iteratively. As illustrated in Figure 1, a proposal is drawn from an ellipse in the source space defined by the current state z_i and a random sample ν from the Gaussian prior. Assuming that the transport map T_θ is continuous in z, it non-linearly maps the ellipse to a continuous, connected curve in the data space that will pass through the current state x_i . The potential function $g(T_\theta(z))$ evaluated in the data space determines if the proposal will be accepted. When the proposal is rejected, the algorithm shrinks the angular bracket around the ellipse, excluding regions that contain the rejected proposal while maintaining support near the current state, and draws a new proposal. As discussed by Murray et al. (2010), this guarantees that the procedure will terminate by accepting a proposal in finite time when the pullback potential function $g \circ T_\theta$ is continuous. However, it excludes potentials that constrain the target distribution to a lower-dimensional manifold in the source space, such as with exact equality constraints. Searching along the ellipse in the source corresponds to gradient-free exploration of the target data space, and the resulting Markov chain in the source $\{z_i\}$ yields the corresponding Markov chain $\{x_i\}$ when passed through the transport map. ESS-Flow has minimal hyperparameters to tune, and one MCMC step is summarized in Algorithm 1.

The convergence of ESS is discussed by Murray et al. (2010) who show that the induced Markov kernel leaves the target distribution invariant. A more detailed convergence analysis is provided by Hasenpflug et al. (2025) and Natarovskii et al. (2021). For completeness we state one of the main results here, adapted to our setting of flow-based models.

Proposition 1 (Geometric convergence of ESS-Flow, Theorem 2.2 by Natarovskii et al. (2021)) Suppose that the pullback potential $z \mapsto g \circ T_{\theta}(z)$ is bounded away from 0 and ∞ on compact sets and has regular tail behavior (details in Natarovskii et al. (2021), Assumption 2.1), then the ESS Markov chain, which we denote $\nu(\cdot, x)$, converges geometrically fast to the target measure π :

```
\|\nu^n(\cdot,z) - \pi\|_{\text{TV}} \le c \left(1 + \|z\|_2\right) \beta^n, \quad \forall z \in \mathbb{R}^d,
```

where $\|\cdot\|_{TV}$ stands for the total variation distance, ν^n is the n-th iteration of the chain, and c>0 and $0<\beta<1$ are constants.

Algorithm 1 ESS-Flow: one MCMC iteration

Require: source $\mathcal{N}(0, I)$, transport map T_{θ} , potential g, current state-potential $(z, g(T_{\theta}(z)))$

```
1: Sample \nu \sim \mathcal{N}(0, I), u \sim U(0, 1)
                                                                             12: function SHRINKBRACKET(\theta, \theta_l, \theta_u)
 2: Sample \theta \sim U(0, 2\pi)
                                                                             13:
                                                                                         if \theta < 0 then
 3: Initialize bracket [\theta_l, \theta_u] \leftarrow [\theta - 2\pi, \theta]
                                                                                               \theta_l \leftarrow \theta
                                                                             14:
 4: while true do
                                                                             15:
                                                                                         else
                                                                                              \theta_u \leftarrow \theta
           z' = z\cos\theta + \nu\sin\theta
 5:
                                                                             16:
           x' = T_{\theta}(z')
 6:
                                                                                         return \theta_l, \theta_u
                                                                             17:
           if \log g(x') > \log g(x) + \log u then
 7:
                 return (z', g(x'))
 8:
 9:
                 \theta_l, \theta_u \leftarrow \text{ShrinkBracket}(\theta, \theta_l, \theta_u)
10:
                 Resample \theta \sim U(\theta_l, \theta_n)
11:
```

4.2 Multi-fidelity sampling with ESS-Flow

Flow-based generative models are defined in continuous time according to equation (1), but in practice the transport map is solved numerically with finite discretization Δ rather than exactly. This means that we in principle have access to a class of approximate priors $\{p_{\theta}^{\Delta}(x):\Delta\in(0,1)\}$ corresponding to different discretization levels, where, intuitively, smaller Δ gives rise to more accurate models.

Our proposed MCMC based sampler could be generalized in different ways to take advantage of such a multi-fidelity setup. The high-level idea is to rely on coarser, and thus computationally cheaper, evaluations of the transport map for a large portion of the evaluations, while ensuring that the final samples nevertheless target a prespecified high-fidelity model. This could be accomplished

by delayed acceptance ESS (Bitterlich et al., 2025), or by methods resembling parallel tempering (Earl & Deem, 2005) or simulated tempering (Marinari & Parisi, 1992) where the annealing temperature is replaced by the discretization level. A simpler approach, which we elaborate here as a proof of concept, is a post-correction of the generated samples using importance weighting.

Let $T_{\theta}^{\Delta}(z)$ denote the transport map with coarse discretization Δ and $T_{\theta}^{\delta}(z)$ with fine discretization $\delta \ll \Delta$. The target distribution from equation (3) for the high-fidelity model can be rewritten as:

$$\pi^{\delta}(z) \propto g(T_{\theta}^{\delta}(z)) \, p(z) = \frac{g(T_{\theta}^{\delta}(z))}{g(T_{\theta}^{\Delta}(z))} g(T_{\theta}^{\Delta}(z)) \, p(z) \propto w \, \pi^{\Delta}(z) \tag{4}$$

We sample from ESS-Flow targeting $\pi^{\Delta}(z)$ using the coarse transport map T^{Δ}_{θ} and re-weight samples with self-normalized importance weights $w_i \propto g(T^{\delta}_{\theta}(z_i))/g(T^{\Delta}_{\theta}(z_i))$ after thinning. The expensive high-fidelity transport map is evaluated only on the final subsampled MCMC samples, not during the ESS itself, significantly reducing computational cost while maintaining accuracy of the target.

5 EXPERIMENTS

We evaluate ESS-Flow on generating materials with target properties and predicting protein backbone structures from sparse inter-residue distances. We compare against state-of-the-art optimization-based methods: D-Flow (Ben-Hamu et al., 2024), PnP-Flow (Martin et al., 2025) and ADP-3D (Levy et al., 2024); and sampling-based method: DAPS (Zhang et al., 2025).

5.1 MATERIAL GENERATION

We use FlowMM (Miller et al., 2024), a flow-based generative model trained with Riemannian flow matching on the MP-20 dataset (Jain et al., 2013), as the prior over materials. A crystalline material c with unit lattice consisting of n atoms is represented by the tuple (a, f, l, β) , where $a \in \{-1, 1\}^{n \times 7}$ are the binarized atomic numbers, $f \in [0, 1]^{n \times 3}$ are the fractional coordinates of each atom, $l \in \mathbb{R}^3_+$ are the lattice cell lengths, and $\beta \in [60^\circ, 120^\circ]^3$ are the lattice angles. We convert the uniform and log-normal source distributions of f, l, β into a standard Gaussian via a change of variables.

To enable comparison with gradient-based methods, we predict material properties using auto-differentiable ALIGNN (Choudhary & DeCost, 2021) models trained on the JARVIS-DFT (Choudhary et al., 2020) dataset, rather than simulation-based procedures. For target properties, we consider bulk modulus, shear modulus, and band gap, choosing target values above the 99th percentile of the prior property distribution. We also evaluate guiding generation toward stable materials, where stability is measured by the predicted energy above hull. Since FlowMM requires specifying the number of atoms per unit lattice, we choose values based on the distribution of atoms for materials with large property values. Target property values, number of atomic sites per lattice, and potential function details are summarized in Table 1.

Table 1: Target properties and values for the material generation experiments

Property P_c	Target y	Atoms n	Potential function $g(c)$	Scale σ_y
Bulk modulus	300 GPa	4	$\exp(-(P_c - y)^2/2\sigma_y^2)$	10 GPa
Shear modulus	200 GPa	4	$\exp(-(P_c - y)^2/2\sigma_y^2)$	10 GPa
Band gap	10 eV	12	$\exp(-(P_c-y)^2/2\sigma_y^2)$	0.1 eV
Energy above hull	-	8	$\exp(-P_c/\sigma_y)$	0.01 eV

We apply dimension-wise learning rates for all methods, as gradient magnitudes vary between (a, f, l, β) . FlowMM outputs $\tilde{a} \in \mathbb{R}^{n \times 7}$, a soft 7-bit encoding that gets rounded to discrete values a. ALIGNN expects integer atomic numbers k_i for the i-th atom to create atom embeddings h_i by selecting the k_i -th row of its embedding matrix E. To maintain differentiability for D-Flow and PnP-Flow, we use the continuous values \tilde{k}_i , derived from \tilde{a}_i , to create soft atom embeddings h_i :

$$z_{ij} = (j - \tilde{k}_i)^2, \quad \boldsymbol{v}_i = \operatorname{softmax}(-\boldsymbol{z}_i/\tau), \quad \boldsymbol{h}_i = \mathbf{E}^\mathsf{T} \boldsymbol{v}_i,$$
 (5)

Table 2: Mean and standard deviation (in parentheses) of absolute errors between the sample properties and targets, and mean and standard deviation of energy above hull values.

Method	Bulk modulus	Shear modulus	Band gap	Energy above hull
Unconditional	198.37 (66.44)	166.71 (22.51)	9.44 (0.98)	2.19 (1.42)
D-Flow	193.86 (69.54)	158.91 (35.13)	9.10 (1.29)	1.52 (1.28)
PnP-Flow	50.98 (31.34)	69.53 (30.83)	5.76 (1.09)	-0.02 (0.06)
DAPS	26.67 (16.83)	55.83 (34.76)	3.34 (1.47)	-0.08 (0.07)
ESS-Flow	9.05 (5.64)	8.30 (6.93)	0.50 (1.09)	-0.20 (0.05)

Table 3: Quality of ESS-Flow samples obtained in different tasks.

Tasks	Validity (%)	Stability (%)	Uniqueness (%)	Novelty (%)	S.U.N. (%)
Bulk modulus	97.3	45.6	37.9	96.2	15.2
Shear modulus	99.8	55.2	18.2	98.9	6.1
Band gap	50.3	29.2	28.9	50.3	21.6
Energy above hull	71.7	53.1	10.0	71.7	6.7

where we set $\tau=0.1$. For DAPS, we avoid this approximation by using Langevin Monte Carlo only for (f,l,β) and sampling a with Metropolis–Hastings using proposals that flip one bit of the binarized atomic numbers. The noising and denoising steps are adapted from the original diffusion formulation to match FlowMM. Hyperparameter details are provided in the Appendix.

We measure performance using the absolute errors between sample properties and targets for conditional generation, and the predicted energy above hull for stability, as shown in Table 2. Even with the continuous approximation for a, D-Flow fails to explore atomic compositions far from initialization. PnP-Flow, which optimizes in data space, is less restricted by this limitation. By avoiding gradient-based exploration of a, DAPS and ESS-Flow perform better, specifically ESS-Flow outperforms all other methods significantly with the lowest errors. The approximation quality is further illustrated in Figure 3, where we see that ESS-Flow can better recover the sharp target distributions.

For ESS-Flow samples, we assess their qualities by computing validity percentage and S.U.N. rate (stability, uniqueness, and novelty). Valid materials have lattice volume greater than 0.1 ų, atomic numbers within range, and balanced charge. After filtering invalid materials, we follow FlowMM and relax lattice structures with CHGNet (Deng et al., 2023), without further density functional theory based relaxations. Stability rate is the proportion of relaxed structures with energy above hull <0.1 eV/atom according to CHGNet. This procedure differs from the predictor used to guide generation of meta-stable materials and can thus produce different results. Uniqueness and novelty rates measure the proportion of unique samples among themselves and with respect to the training set, respectively, using pymatgen's StructureMatcher (Ong et al., 2013) with default settings. Results computed over 1000 generated samples are reported in Table 3.

Validity and stability rates are limited by the FlowMM prior, which may produce unrelaxed structures and has compositional validity of approximately 83%. The uniqueness rate could be improved by using parallel MCMC chains instead of one chain with thinning, which may result in correlated samples. The S.U.N. rates are naturally low compared to unconditional generation, but they should be viewed in light of the fact that we are (successfully) targeting extreme values for the desired properties, with target values set to the 99th percentile based on the unconditional model. Despite this challenging controlled generation problem, the resulting S.U.N. rates indicate that ESS-Flow does generate a subset of materials that are good candidates for further exploration.

5.1.1 EVALUATION OF MULTI-FIDELITY SAMPLING

We perform preliminary evaluation of the suggested proof of concept for multi-fidelity sampling in Section 4.2. The ESS-Flow samples obtained from evaluating the potential with a coarse discretization $\Delta=1/50$, is re-weighted with importance weights by using much finer discretization for the ODE $\delta=1/1000$ as described in equation (4). The weighted samples have a reasonable effective sample size of 68.1% and 74.1% for the bulk modulus and shear modulus tasks respectively.

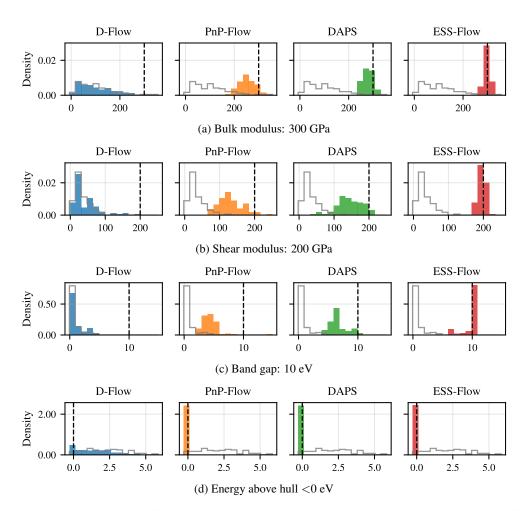


Figure 3: Sample property distributions. Property values are along the X-axis, prior property distribution is shown in gray, and target values are shown as dotted lines. ESS-Flow samples are near target values across all tasks.

However, in the band gap and stability tasks that have sharper target distributions, they are lower, 6.8% and 11.7% respectively. This is a shortcoming of the simple importance re-weighting approach which assigns large weights for samples with low potentials under the coarse discretization.

5.2 PROTEIN STRUCTURE PREDICTION

For protein backbone structure prediction from partial inter-residue distances, we follow Levy et al. (2024). We use Chroma (Ingraham et al., 2023), a pretrained diffusion model, as our prior over protein backbone structures denoted by $x \in \mathbb{R}^{4n \times 3}$, which are the 3D coordinates of all four heavy atoms from n amino acid residues. Since we require a deterministic mapping between source and data samples, we modify Chroma's random protein graph construction to use k-nearest neighbors and generate samples with the probability flow ODE.

Levy et al. (2024) compute all pairwise distances between α -carbons in protein PDB: 7r5b (Warstat et al., 2023) containing 147 residues and randomly sample distances without noise. However, they note that this is simplified, as nuclear magnetic resonance spectroscopy cannot probe distances above 6 Å. We account for this by sampling only 300 out of 330 distances that are <6 Å and adding Gaussian noise $\mathcal{N}(0,0.5^2)$, as the protein is deposited at the resolution of 1.77 Å. This results in a highly underdetermined inverse problem, suggesting that a Bayesian approach is suitable for capturing the diversity of structures that agree with the observed data. Conditioned on these observations, we generate 10 backbone structures with D-Flow, ADP-3D, DAPS, and ESS-Flow. We report the

 L^2 distance d_y between observations and corresponding sample distances, and the root mean square distance RMSD_{gt} between the ground truth structure and samples in Table 4.

Table 4: Mean and standard deviation (in parenthesis) of the metrics for the sampled structures.

Method	d_y	$RMSD_{gt}$	min. RMSD _{gt}	ELBO
Unconditional	80.21 (9.69)	16.98 (0.83)	15.53	8.70 (0.20)
D-Flow	46.54 (8.58)	14.44 (1.35)	11.54	8.64 (0.35)
ADP-3D	3.43 (0.34)	11.45 (1.52)	8.91	-5.68 (1.24)
DAPS	11.79 (1.67)	11.41 (1.17)	9.47	-8.07 (1.27)
ESS-Flow	37.02 (5.06)	13.55 (1.32)	10.63	8.89 (0.21)

Following Levy et al. (2024), we use the lower bound on the log marginal data likelihood (ELBO) from Chroma as a measure of structural realism. While L^2 distances suggest that ADP-3D and DAPS samples best match observations, the ELBO values and Figure 4 reveal that resulting structures are highly unnatural. In both methods, as noise levels are annealed, regularization from the pretrained generative model diminishes, shifting optimization toward maximum likelihood estimates. This results in prioritizing a good data fit (low RMSD) at the expense of weaker prior regularization (unrealistic structures). ESS-Flow explicitly enforces the prior, resulting in more realistic samples. We observe a limitation, however, where ESS-Flow does not explore the target distribution as effectively as in the material generation experiments and tends to become trapped in modes. Therefore, samples are obtained from 10 independent parallel MCMC chains.

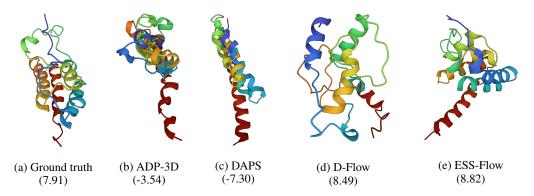


Figure 4: Ground truth protein structure PDB: 7r5b, conditional sample with the lowest RMSDgt from each method, and their ELBO from Chroma (in parenthesis). While samples from ADP-3D and DAPS resemble partially collapsed polymers, ESS-Flow achieves a better trade-off between data fidelity and sample realism.

6 Conclusion

We introduce ESS-Flow, a training-free sampling method for controlled generation with pretrained flow-based generative models. The method requires minimal hyperparameter tuning and is entirely gradient-free, making it particularly valuable for problems where gradient-based optimization struggles, such as when the generative process involves quantization, or evaluating the potential requires non-differentiable simulations. However, this gradient-free approach limits ESS-Flow's effectiveness when the prior does not well inform the target distribution, such as in noiseless image inpainting tasks. Currently, we use moderate numbers of function evaluations in the ODE solver, fewer than what is typically used for unconditional generation. We propose and evaluate a multi-fidelity setup where samples obtained by using coarse ODE discretization are re-weighted with importance weights, as a proof of concept. Future work could explore more principled combinations of coarse and fine ODE discretization to better guide exploration of target distributions. Additionally, developing adaptive strategies for problems where the prior poorly covers the target distribution could expand the applicability of the method to a broader range of controlled generation tasks.

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.
- Rebecca F Alford, Andrew Leaver-Fay, Jeliazko R Jeliazkov, Matthew J O'Meara, Frank P DiMaio, Hahnbeom Park, Maxim V Shapovalov, P Douglas Renfrew, Vikram K Mulligan, Kalli Kappel, et al. The rosetta all-atom energy function for macromolecular modeling and design. *Journal of chemical theory and computation*, 13(6):3031–3048, 2017.
- Amr Alhossary, Stephanus Daniel Handoko, Yuguang Mu, and Chee-Keong Kwoh. Fast, accurate, and reliable molecular docking with quickvina 2. *Bioinformatics*, 31(13):2214–2216, 2015.
- Heli Ben-Hamu, Omri Puny, Itai Gat, Brian Karrer, Uriel Singer, and Yaron Lipman. D-Flow: Differentiating through flows for controlled generation. In *International Conference on Machine Learning*, pp. 3462–3483. PMLR, 2024.
- Kevin Bitterlich, Daniel Rudolf, and Björn Sprungk. Delayed acceptance elliptical slice sampling. *PAMM*, 25(1):e70005, 2025.
- Gabriel Victorino Cardoso, Yazid Janati El Idrissi, Sylvain Le Corff, and Eric Moulines. Monte carlo guided diffusion for bayesian linear inverse problems. In *ICLR International Conference on Learning Representations*, 2024.
- Kamal Choudhary and Brian DeCost. Atomistic line graph neural network for improved materials property predictions. *npj Computational Materials*, 7(1):185, 2021.
- Kamal Choudhary, Kevin F Garrity, Andrew CE Reid, Brian DeCost, Adam J Biacchi, Angela R Hight Walker, Zachary Trautt, Jason Hattrick-Simpers, A Gilad Kusne, Andrea Centrone, et al. The joint automated repository for various integrated simulations (JARVIS) for data-driven materials design. *npj computational materials*, 6(1):173, 2020.
- Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *International Conference on Learning Representations*, 2023.
- Hyungjin Chung, Jeongsol Kim, and Jong Chul Ye. Diffusion models for inverse problems. *arXiv* preprint arXiv:2508.01975, 2025.
- Giannis Daras, Hyungjin Chung, Chieh-Hsin Lai, Yuki Mitsufuji, Jong Chul Ye, Peyman Milanfar, Alexandros G Dimakis, and Mauricio Delbracio. A survey on diffusion models for inverse problems. *arXiv preprint arXiv:2410.00083*, 2024.
- Bowen Deng, Peichen Zhong, KyuJung Jun, Janosh Riebesell, Kevin Han, Christopher J Bartel, and Gerbrand Ceder. CHGNet as a pretrained universal neural network potential for charge-informed atomistic modelling. *Nature Machine Intelligence*, 5(9):1031–1041, 2023.
- Prafulla Dhariwal and Alexander Nichol. Diffusion models beat GANs on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- David J Earl and Michael W Deem. Parallel tempering: Theory, applications, and new perspectives. *Physical Chemistry Chemical Physics*, 7(23):3910–3916, 2005.
- Matthew Graham and Amos Storkey. Asymptotically exact inference in differentiable generative models. In *Artificial Intelligence and Statistics*, pp. 499–508. PMLR, 2017.
- Mareike Hasenpflug, Viacheslav Telezhnikov, and Daniel Rudolf. Reversibility of elliptical slice sampling revisited. *Bernoulli*, 31(2):1377–1401, 2025.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- John B Ingraham, Max Baranov, Zak Costello, Karl W Barber, Wujie Wang, Ahmed Ismail, Vincent Frappier, Dana M Lord, Christopher Ng-Thow-Hing, Erik R Van Vlack, et al. Illuminating protein space with a programmable generative model. *Nature*, 623(7989):1070–1078, 2023.

- Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, et al. Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL materials*, 1(1), 2013.
 - Yazid Janati, Badr Moufad, Alain Durmus, Eric Moulines, and Jimmy Olsson. Divide-and-conquer posterior sampling for denoising diffusion priors. *Advances in Neural Information Processing Systems*, 37:97408–97444, 2024.
 - Filip Ekström Kelvinius, Zheng Zhao, and Fredrik Lindsten. Solving linear-gaussian bayesian inverse problems with decoupled diffusion sequential monte carlo. In *International Conference on Machine Learning*, 2025.
 - Jeongsol Kim, Bryan Sangwoo Kim, and Jong Chul Ye. Flowdps: Flow-driven posterior sampling for inverse problems. *arXiv preprint arXiv:2503.08136*, 2025.
 - Axel Levy, Eric R Chan, Sara Fridovich-Keil, Frédéric Poitevin, Ellen D Zhong, and Gordon Wetzstein. Solving inverse problems in protein space using diffusion-based priors. *arXiv* preprint *arXiv*:2406.04239, 2024.
 - Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023.
 - Xingchao Liu, Chengyue Gong, et al. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *International Conference on Learning Representations*, 2023.
 - Enzo Marinari and Giorgio Parisi. Simulated tempering: a new monte carlo scheme. *Europhysics letters*, 19(6):451, 1992.
 - Ségolène Tiffany Martin, Anne Gagneux, Paul Hagemann, and Gabriele Steidl. PnP-Flow: Plugand-play image restoration with flow matching. In *International Conference on Learning Representations*, 2025.
 - Benjamin Kurt Miller, Ricky TQ Chen, Anuroop Sriram, and Brandon M Wood. FlowMM: Generating materials with Riemannian flow matching. In *International Conference on Machine Learning*, pp. 35664–35686. PMLR, 2024.
 - Iain Murray, Ryan Adams, and David MacKay. Elliptical slice sampling. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 541–548. JMLR Workshop and Conference Proceedings, 2010.
 - Viacheslav Natarovskii, Daniel Rudolf, and Björn Sprungk. Geometric convergence of elliptical slice sampling. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, pp. 7969–7978. PMLR, 18–24 Jul 2021.
 - Shyue Ping Ong, William Davidson Richards, Anubhav Jain, Geoffroy Hautier, Michael Kocher, Shreyas Cholia, Dan Gunter, Vincent L Chevrier, Kristin A Persson, and Gerbrand Ceder. Python materials genomics (pymatgen): A robust, open-source python library for materials analysis. *Computational Materials Science*, 68:314–319, 2013.
 - Ashwini Pokle, Matthew J Muckley, Ricky TQ Chen, and Brian Karrer. Training-free linear image inverses via flows. *Transactions on Machine Learning Research*, 2024.
 - Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman, and Ricky TQ Chen. Multisample flow matching: Straightening flows with minibatch couplings. In *International Conference on Machine Learning*, pp. 28100–28127. PMLR, 2023.
 - Vishal Purohit, Matthew Repasky, Jianfeng Lu, Qiang Qiu, Yao Xie, and Xiuyuan Cheng. Consistency posterior sampling for diverse image synthesis. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 28327–28336, 2025.
 - Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.

- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*, pp. 1–34, 2024.
- Zifan Wang, Alice Harting, Matthieu Barreau, Michael M Zavlanos, and Karl H Johansson. Source-guided flow matching. *arXiv preprint arXiv:2508.14807*, 2025.
- Robin Warstat, Mehrosh Pervaiz, Pierre Regenass, Marius Amann, Karin Schmidtkunz, Oliver Einsle, Manfred Jung, Bernhard Breit, Martin Hügle, and Stefan Günther. A novel pan-selective bromodomain inhibitor for epigenetic drug design. *European Journal of Medicinal Chemistry*, 249:115139, 2023.
- Luhuan Wu, Brian Trippe, Christian Naesseth, David Blei, and John P Cunningham. Practical and asymptotically exact conditional sampling in diffusion models. *Advances in Neural Information Processing Systems*, 36:31372–31403, 2023.
- Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song. Improving diffusion inverse problem solving with decoupled noise annealing. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 20895–20905, 2025.
- Zheng Zhao, Ziwei Luo, Jens Sjölund, and Thomas B. Schön. Conditional sampling within generative diffusion models. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 383(2299):20240329, 2025.

A APPENDIX

A.1 ADDITIONAL DETAILS FOR THE EXPERIMENTS

The hyperparameters for methods used in the material generation experiment and protein structure prediction are given in Table 5 and 6 respectively.

Table 5: Hyperparameters for the material generation tasks

Method / Parameter	Bulk modulus	Shear modulus	Band gap	Energy above hull
Unconditional				
ODE integration steps N	50	50	50	50
D-Flow				
ODE integration steps N	50	50	50	50
Optimizer	SGD	SGD	SGD	SGD
Optimization steps N_{opt}	100	100	100	100
Learning rate $(\boldsymbol{a}, \boldsymbol{l}, \boldsymbol{\beta})$	10^{-6}	10^{-5}	10^{-4}	10^{-4}
Learning rate (f)	10^{-8}	10^{-7}	10^{-6}	10^{-6}
PnP-Flow				
Optimizer	SGD	SGD	SGD	SGD
Optimization steps N_{opt}	100	100	100	100
Learning rate decay	linear	linear	linear	linear
Min. learning rate	0	0	0	0
Max. learning rate (a)	10^{-4}	10^{-3}	10^{-1}	10^{-2}
Max. learning rate (f)	10^{-6}	10^{-5}	10^{-3}	10^{-4}
Max. learning rate $(\boldsymbol{l}, \boldsymbol{\beta})$	10^{-5}	10^{-4}	10^{-2}	10^{-3}
ODE integration steps N	100 - n	100 - n	100 - n	100 - n
DAPS				
ODE integration steps N	5	5	5	5
Optimization steps N_{opt}	100	100	100	100
Sampling steps N_s	20	20	20	20
Learning rate decay	linear	linear	linear	linear
Min. learning rate	$10^{-2}\alpha$	$10^{-2}\alpha$	$10^{-2}\alpha$	$10^{-2}\alpha$
Max. learning rate (α_f)	10^{-3}	10^{-3}	10^{-5}	10^{-4}
Max. learning rate $(\alpha_{l}, \alpha_{\beta})$	10^{-2}	10^{-2}	10^{-4}	10^{-3}
$p(x_0 x_t)$ scale $(r_{\boldsymbol{a}})$	3(1-t)	3(1-t)	3(1-t)	3(1-t)
$p(x_0 x_t)$ scale (r_f, r_l, r_{beta})	1-t	1-t	1-t	1-t
ESS-Flow				
ODE integration steps N	50	50	50	50
MCMC steps	1200	1200	1200	1200
Burn-in	200	200	200	200
Thinning factor	10	10	10	10

Table 6: Hyperparameters for the protein structure prediction

Method / Parameter	
Unconditional	
ODE integration steps N	20
D-Flow	
ODE integration steps N	20
Optimizer	SGD
Optimization steps N_{opt}	500
Learning rate schedule	CyclicLR
Min. learning rate	10^{-6}
Max. learning rate	10^{-5}
Step size up/down	125
Momentum	0.9
ADP-3D	
Optimizer	SGD
Optimization steps N_{opt}	1000
Learning rate	0.67
Momentum	0.99
DAPS	
ODE integration steps N	5
Optimization steps N_{opt}	1000
Sampling steps N_s	5
Learning rate decay	linear
Min. learning rate	10^{-5}
Max. learning rate	10^{-4}
$p(x_0 x_t)$ scale	2t
ESS-Flow	
ODE integration steps N	20
MCMC steps	201
Burn-in	200