LEARNING TO SAMPLE FROM MULTIMODAL DISTRIBUTIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Markov chain Monte Carlo (MCMC) is an established method for sampling from complex probability distributions by building Markov chains whose state transitions occur by a proposal-acceptance mechanism. MCMC however tends to fail on distributions whose modes are separated by wide low-probability regions. Neural networks have been integrated into MCMC to try and address this challenge, but existing approaches cannot guarantee that the resulting chain asymptotically produces samples from the target distribution, or even that they cover all modes. In this work, we use constrained optimization to put forward a new way of training neural MCMC that reduces samples correlation, ensures the coverage of all modes, and lower bounds the proposal acceptance rate. We prove that, under mild conditions, any solution of this constrained problem results in an ergodic Markov chain whose empirical measure converges to the target distribution. We then propose an effective algorithm to tackle this problem and illustrate its performance on a variety of applications from mixture models to physical systems, and reliability engineering showing that it matches or outperforms existing baselines.

1 INTRODUCTION

Sampling from probability distributions is a fundamental problem in statistics (Robert & Casella, 2004; McElreath, 2018), machine learning (Koller & Friedman, 2009; Andrieu et al., 2003; Du & Mordatch, 2019), and the natural sciences (Foreman et al., 2021; Nowak et al., 2010; Schneider et al., 2024; Vrugt, 2016). The problem is challenging because the distributions we wish to sample from (*target distributions*) are only known in an unnormalized form. This is the case, for instance, in Bayesian inference, where the posterior distribution over model parameters is only obtained up to a normalization factor (evidence). Markov chain Monte Carlo (MCMC) methods (Robert & Casella, 2004) are ubiquitous in this setting, because they yield, under mild conditions, asymptotically unbiased samples from the target distribution by simulating transitions of a Markov chain such that the average time spent in each state is proportional to its probability under the target distribution. Popular choices include variants of the discretized Langevin (Roberts & Tweedie, 1996b) and Hamiltonian dynamics (Duane et al., 1987; Homan & Gelman, 2014; Hoffman & Sountsov, 2022), which have provably fast convergence rates for log-concave target distributions (Roberts & Tweedie, 1996b; Livingstone et al., 2019). However, they are also known to not work well for multimodal target distributions, i.e., distributions with multiple modes separated by regions of low probability, that render leaving any mode using small *incremental* moves unlikely (Łatuszyński et al., 2026).

Neural networks (NNs) have been proposed to mitigate this challenge by providing a rich and flexible parametrizations of the state transitions of a Markov chain that allow to *directly* transition between the states of multiple modes (Song et al., 2017; Levy et al., 2018; Hoffman et al., 2019; Spanbauer et al., 2020). Yet, not every repeated application of a NN yields an occupancy distribution of states that is proportional to the target distribution. Training on decorrelating successive samples as in Pasarica & Gelman (2010); Levy et al. (2018) or adversarially making them indistinguishable from a bootstrapped target as in Song et al. (2017); Spanbauer et al. (2020); Egorov et al. (2024) does not guarantee that the target distribution is eventually reached. This comes from the technical fact that these approaches do not guarantee *ergodicity*. For this property to hold, it is *necessary* (but *not sufficient*) that the target distribution needs to be a *fixed point* of the Markov chain's transition dynamics which can be realized through self-inverse neural networks (Neklyudov et al., 2020). This means that the occupancy distribution of states remains unchanged in expectation

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

once it is proportional to the target distribution. It remains an open question how to design training objectives that are also *sufficient* to guarantee that this stationary occupancy distributions is *unique* and *reachable* from any initial state (Cannella & Tarokh, 2022; Egorov et al., 2024). In other words, that the resulting Markov chain is ergodic and thus asymptotically produces unbiased samples from the target distribution.

This work tackles this challenge by casting the problem of training neural MCMC methods as a constrained learning problem: we maximize the decorrelation of consecutive samples subject to three constraints on the transition kernel that enforce desirable properties of an ergodic Markov chain, namely (i) coverage of all modes, (ii) fidelity to the target within each mode, and (iii) high acceptance probability. We prove that any kernel satisfying theses constraints yields a uniformly ergodic Markov chain under mild assumptions, with a convergence rate explicit in the constraint thresholds. We then use dual learning methods (Chamon & Ribeiro, 2020; Chamon et al., 2023) to develop an algortihm for training these kernels. We show the effectiveness of our method accross applications, including mixture models, physical systems, reliability engineering, and arbitrary conditioning of probabilistic queries to (trained) generative models.

2 Background

This work assumes that the target distribution p over a measurable space $(\mathcal{X}, \mathcal{B})$ where $\mathcal{X} \subseteq \mathbb{R}^d$ and $\mathcal{B}$ is a σ-algebra over $\mathcal{X}$. By a slight abuse of notation, we will denote the target distribution by its density $p\colon \mathcal{X} \to [0, \infty)$, which is known up to a normalizing constant Z, i.e. $p(x) = Z \cdot \tilde{p}(x)$.

2.1 Markov Chain Monte Carlo (MCMC)

A Markov chain simulates a correlated sequence of *states* $x_0, \dots, x_N \in \mathcal{X}$ obtained by repeatedly applying a *transition kernel* $t_\theta(x'|x)$ with parameters θ that determines the next state x' from the current state x (Robert & Casella, 2004). To obtain (unbiased) samples from a target distribution, the average time spent in state x must be proportional to $p(x)$. For this property to hold, it is *necessary* (Neklyudov et al., 2020) the target distribution to be a *fix point* of the kernel t_θ, i.e.,

$$\int_{\mathcal{X}} t_\theta(x' \mid x)\, p(x)\, dx = p(x'), \qquad \forall\, x' \in \mathcal{X}. \tag{1}$$

Classically, this is achieved by combining a proposal distribution $q_\theta(x' \mid x)$, from which a candidate for the next state x' is drawn given the current state x, with an acceptance test. If a proposed state x' is accepted it becomes the next state of the Markov chain. Otherwise, the next state remains x. Choices of proposals range from Gaussian (Metropolis et al., 1953) to gradient-informed discretizations of Langevin (Roberts & Tweedie, 1996b) and Hamiltonian dynamics (Duane et al., 1987; Hoffman & Sountsov, 2022). When this fixed point is *unique* and *reachable* from any initial state x_0, the Markov chain is said to be *ergodic* (Meyn & Tweedie, 1993; Neklyudov et al., 2020). For multimodal target distributions, a local moves such as the classical $x' \sim \mathcal{N}(x, \theta \cdot I)$ with small θ (Metropolis et al., 1953) are unlikely to cross regions of low probability and reach other modes. Instead, proposals $x' = f_\theta(x)$ can be generated by a NN f_θ that is trained to *directly* transition between the modes of the target distribution.

2.2 Neural Proposals

This idea of neural proposals has been unified in the Involutive MCMC (iMCMC) algorithm (Neklyudov et al., 2020). It operates in an *augmented state space* $\mathcal{X} \times \mathcal{V}$ by introducing an *auxiliary variable* $v \in \mathcal{V}$ of *known* distribution $p_v(v)$ and letting $p(x, v) = p(x)\, p_v(v)$. Follwing neural proposals in this augmented space causes next state pairs to be

$$(x', v') = \begin{cases} f_\theta(x, v), & \text{with probability } \alpha(x, v)\,, \\ (x, v), & \text{otherwise,} \end{cases} \tag{2}$$

and the transition density thus becomes a mixture of two Dirac distributions $\delta(\cdot)$, as

$$t_\theta(x', v' \mid x, v) = \delta\left(\begin{pmatrix} x' \\ v' \end{pmatrix} - f_\theta\,(x, v)\right)\, \alpha(x, v) + \delta\left(\begin{pmatrix} x' \\ v' \end{pmatrix} - \begin{pmatrix} x \\ v \end{pmatrix}\right)\, \big(1 - \alpha(x, v)\big), \tag{3}$$

one at the proposed state $f_\theta(x, v)$ and one at the current state (x, v), weighted by the Metropolis–Hastings acceptance probability of the proposed state,

$$\alpha(x, v) = \min\left\{1,\ \frac{p(f(x, v))}{p(x, v)}\ \left|\frac{\partial f}{\partial(x, v)}\right|\right\}, \tag{4}$$

where $\left|\frac{\partial f}{\partial(x,v)}\right|$ denotes the absolute Jacobian determinant of f. Restricting f_θ to be an *involution*, $f_\theta = f_\theta^{-1}$, provably (Neklyudov et al., 2020) fulfills Eq. (1) and can be enforced by using specialized NN architectures (Levy et al., 2018; Hoffman et al., 2019; Spanbauer et al., 2020). We can interleave applications of f_θ with resampling of the auxiliary variable $v \sim p_v$ similar to Gibbs sampling (Geman & Geman, 1984) to form a Markov chain.

Training. While classical MCMC often produces a sequence of small local moves and highly correlated samples, a neural proposal can be trained to produce successive states that are as uncorrelated as possible (Levy et al., 2018). The standard proxy for a small *lag-one autocorrelation* is the *expected squared jump distance* (ESJD) (Pasarica & Gelman, 2010) which is for an involutive kernel defined in terms of Eq. (3) as

$$\mathrm{ESJD}(\theta) = \mathbb{E}_{x\sim p,\, v\sim p_v}\Big[\alpha(x, v)\, \|x' - x\|_2^2\Big], \qquad \text{for} (x', v') = f_\theta(x, v). \tag{5}$$

Alternative training objectives are discussed in Section 5. Yet, a large jump distance alone does not guarantee that the chain visits every mode of the target, nor that it converges at all. Next we show how to use constrained optimiztation to ensure that it does.

3 Training Neural Proposals with Constrained Learning

In this section, we describe a constrained learning strategy for training neural proposals f_θ that jointly and explicitly enforces (i) adaption to the target distribution within each mode (Section 3.1), (ii) mixing between modes (Section 3.2), and (iii) acceptance of proposed states (Section 3.3). We later show in Section 3.4 that such a proposal induces an ergodic Markov chain, i.e. one that is guaranteed to eventually produce samples from the target distribution.

3.1 Adapting the chain to the shape of the target distribution

Ideally, the proposal should be as close as possible to the target distribution. To ensure this is the case, we enforce that the Fisher divergence between $q_\theta(\cdot|\hat{x})$, the proposal distribution at state $\hat{x}$, and the target distribution p is small. Explicitly, for $\epsilon > 0$, we require

$$\mathrm{F}\left(q_\theta(\cdot \mid \hat{x}) \| p\right) = \mathbb{E}_{x\sim q_\theta(\cdot|\hat{x})}\Big[\, \|\nabla_x \log q_\theta(x|\hat{x}) - \nabla_x \log \tilde{p}(x)\|_2^2 \,\Big] \leq \epsilon. \tag{6}$$

Here, we choose the *reverse* divergence because it only requires samples from $q_\theta(\cdot|\hat{x})$. By using an involutive NN f_θ, the marginal proposal distribution

$$q_\theta(x'|x) = \int \delta\left(f_\theta(x, v) - (x', v')^T\right) \left|\frac{\partial f}{\partial(x, v)}\right| p_v(v|x)\ dv\, dv' \tag{7}$$

is tractable and straightforward to sample (see Appendix A.1). In addition, the Fisher divergence only depends on the score $\nabla_x \log p(x) = \nabla_x \log \tilde{p}(x)$ of the target distribution and is therefore independent of the normalizing constant Z. This eases the choice of ϵ in practice when Z is unknown.

However, choosing the reverse direction of divergence measures comes at a cost: it is *mode-seeking*, i.e., the distance between the distributions $q_\theta(\cdot|\hat{x})$ and p can be at an negligible level while missing isolated components of the target distribution (Papamakarios et al., 2021; Wenliang, 2020). We therefore consider an additional condition to mitigate this effect.

3.2 Ensure mixing between the modes

To ensure that the Markov chain transitions not only within but also between modes, we should use a *forward* divergence measure, which is *mode-covering* (Papamakarios et al., 2021). Nevertheless,

these are intractable in this settting sicne they require integrating w.r.t. the target distribution p which is only known up to a normalizing constant Z. Hence, we instead consider a *reference distribution* $\hat{p}$ that is easy to sample from and whose support contains the support of the target distribution p and ask that the proposal distribution q_θ to be close to this reference, namely that for $\kappa > 0$,

$$\mathrm{KL}\left(\hat{p} \| q_\theta(\cdot|\hat{x})\right) = \int_{\mathcal{X}} \hat{p}(x) \log \frac{\hat{p}(x)}{q_\theta(x|\hat{x})} dx = \mathbb{E}_{x\sim\hat{p}}\Big[\log \hat{p}(x) - \log q_\theta(x|\hat{x})\Big] \leq \kappa. \tag{8}$$

Since both $\hat{p}$ and q_θ are (properly normalized) probability distributions, the Kullback-Leibler divergence Eq. (8) is non-negative and zero if and only if the two distributions are identical. Thus, choosing $\kappa \approx 0$ would end up distilling the reference distribution into the neural proposal, which is *not* not the desired behavior.

Instead, Eq. (8) is meant to ensure the coverage of the support of $\hat{p}$. For that to happen, it is enough that the KL divergence is finite. Indeed, note that the KL divergence is finite if and only if $\hat{p} \ll q_\theta(\cdot|\hat{x})$. Hence by taking a larger value of κ, we can ensure that all modes are visited without substantially pushing the proposal distribution $q_\theta(\cdot|\hat{x})$ away from the target distribution p even if the reference distribution $\hat{p}$ is not a good approximation of the target distribution p.

How to choose the reference? If sufficient domain knowledge is not available, we can construct a reference distribution from a set of support points $\mathcal{S} = \{\hat{x}_i\}_{i=1}^N$ that cover the support of the target distribution. Simple gradient-based optimization can be used to find local maxima of the target distribution from diverse initializations as in Pompe et al. (2020); Noble et al. (2024): Starting from a Sobol sequence (Sobol', 1967) of points we build a graph to find candidates that are local maxima and refine them with L-BFGS (Liu & Nocedal, 1989) to obtain the support points $\mathcal{S}$ (see Appendix A.2 for details). These points, can then be used to construct a *mixture of Laplace approximation* of the target distribution without access to the normalization constant Z. Explicitly,

$$\hat{p}(x) = \sum_i \omega_i \, \mathcal{N}(x; \hat{x}_i, \Sigma_i), \quad \Sigma_i = \left(-\nabla^2 \log p(\hat{x}_i)\right)^{-1}, \text{and} \quad \omega_i = \frac{p(\hat{x}_i)\det(2\pi\Sigma_i)^{1/2}}{\sum_j p(\hat{x}_j)\det(2\pi\Sigma_j)^{1/2}}. \tag{9}$$

3.3 Improving acceptance

The two conditions so far control where the proposal places its mass, but a proposed state only becomes the next state of the chain if it passes the acceptance test in (4). The acceptance probability is not determined by the quality of the proposed state alone, i.e. by whether it lands in a region of high target density, but also by how likely the move is to be reversed. For an involutive map, this is controled by the Jacobian determinant in (4). A kernel that satisfies the conditions (6) and (8) could therefore still have almost every proposal rejected. To ensure this does not happen, we therefore impose a lower bound $\delta > 0$ on the acceptance probability as in

$$\mathbb{E}_{v\sim p_v}\big[\alpha(\hat{x}, v)\big] = \mathbb{E}_{v\sim p_v}\left[\min\left\{1, \; \frac{p(f_\theta(\hat{x}, v))}{p(\hat{x}, v)}\left|\frac{\partial f_\theta}{\partial(\hat{x}, v)}\right|\right\}\right] \geq \delta. \tag{10}$$

In words, this lower bounds the average acceptance probability of the proposal at $\hat{x}$. During training we replace the minimum in Eq. (10) by a smooth surrogate (see Appendix A.1).

3.4 The Final Training Problem

We combine the conditions (6), (8) and (10) with the training objective in Eq. (5) to obtain the final optimization problem used to train our neural proposals, namely,

$$\begin{aligned} \underset{\theta}{\text{maximize}} \quad & \mathrm{ESJD}(\theta) \\ \text{subject to} \quad & \mathrm{F}\left(q_\theta(\cdot|\hat{x}_i)\|p\right) \leq \epsilon, \quad \forall \hat{x}_i \in \mathcal{S}, \\ & \mathrm{KL}\left(\hat{p}\|q_\theta(\cdot|\hat{x}_i)\right) \leq \kappa, \quad \forall \hat{x}_i \in \mathcal{S}, \\ & \mathbb{E}_{v\sim p_v}\big[\alpha(\hat{x}_i, v)\big] \geq \delta, \quad \forall \hat{x}_i \in \mathcal{S} \end{aligned} \tag{11}$$

Recall that the Fisher constraint ϵ matches the proposal shape to the target within each mode, the coverage constraint κ keeps the proposal mass spread across the modes, and the acceptance constraint δ keeps the chain moving. The objective favors the decorrelation of consecutive samples.

Though the constraints in (11) may appear arbitrary, they are not only desirable properties of the proposal. In fact, they are sufficient to ensure that a proposal induces, under mild assumptions, an ergodic Markov chain. This fact is established next.

Let lowercase roman symbols be densities p, $\hat{p}$, $q_\theta(\cdot|x)$ and t_θ, and fraktur symbols $\mathfrak{p}$, $\widehat{\mathfrak{p}}$, $\mathfrak{q}_\theta(\cdot|x)$ and $\mathfrak{t}_\theta$ denote the measures they induce. Let further $\|\mu - \nu\|_{\mathrm{TV}} := \sup_B |\mu(B) - \nu(B)|$ denote the total variation distance between two measures μ and ν and consider the following assumptions:

- **A1 (Architecture)** The proposal density $q_\theta(\cdot|x) \in C^1(\mathcal{X})$ is smooth and positive.
- **A2 (Locally log-concave target)** There exist disjoint, open, and convex sets $U_1, \ldots, U_M$ for which the target distribution $p \in C^2(\mathcal{X})$ is smooth and positive and $-\nabla^2 \log p(x) \succeq mI$ for all $x \in U_j$, $j = 1, \ldots, M$.
- **A3 (Reference quality)** The reference distribution is such that $\sum_{c=0}^{M} \left|\widehat{\mathfrak{p}}(U_c) - \mathfrak{p}(U_c)\right| \leq \tau$.
- **A4 (Generalization)** The constraints in (11) hold for every $x \in \mathcal{X}$, i.e.

$$\mathrm{KL}\left(\hat{p} \| q_\theta(\cdot|x)\right) \leq \kappa, \quad \mathrm{F}\left(q_\theta(\cdot|x) \| p\right) \leq \epsilon, \text{ and } \quad \mathbb{E}_{v \sim p_v}\left[\alpha(x, v)\right] \geq \delta.$$

Here, we define $U_0 := \mathcal{X} \setminus \bigcup_{j=1}^{M} U_j$ and $\eta := \mathfrak{p}(U_0)$ as the residual mass. Observe that A1 holds for the architecture introduced in Appendix A.1. A2 imposes *no* global log-concavity, only on the sets U_j, e.g. balls around the modes of a well-separated mixture. A3 captures the quality of the reference distribution. Our experiments (Section 6) suggest that its quality is not critical, as long as it covers the modes. Future work may tighten the theoretical analysis in this regard as the current proof uses (loose) upper bounds. A4 can be influenced by training: tighter thresholds κ, ϵ, δ can guarantee a (better) convergence rate, while making it harder to satisfy the constraints during training.

With these assumptions in place, the main theorem guarantees that our training method yields an ergodic Markov chain, for any feasible solution of the constrained program (11).

Theorem 1. *Let Assumptions A1–A4 hold and set*

$$\rho(\epsilon, \delta, \kappa) \;:=\; 2(1-\delta) \;+\; \sqrt{2\kappa} + \tau \;+\; \sqrt{\epsilon/m} \;+\; 2\eta. \tag{12}$$

If $\rho < 1$, then the CI-MCMC kernel is uniformly ergodic, i.e.

$$\|\mathfrak{t}_\theta^n(x, \cdot) - \mathfrak{p}\|_{\mathrm{TV}} \;\leq\; \tfrac{1}{2}\,\rho(\epsilon, \delta, \kappa)^n \qquad \textit{for all } x \in \mathbb{R}^d,\; n \geq 1. \tag{13}$$

A proof is provided in Appendix C. Each threshold in the problem contributes one term of ρ: the geometric convergence rate becomes faster as $\kappa \to 0$, $\epsilon \to 0$, $\delta \to 1$, and $\tau, \eta \to 0$.

4 Algorithm

In the following we describe how to algorithmically obtain solutions to the constrained optimization problem. We call the resulting method *Constrained Involutive MCMC* (CI-MCMC).

Training. We first locate the modes of the target and construct the reference distribution $\hat{p}$ of Eq. (9) from them; the search procedure is given in Appendix A.2. We then solve the constrained program by primal-dual optimization, applying the constrained learning framework of Chamon & Ribeiro (2020); Chamon et al. (2023). Every inequality of the program receives a non-negative multiplier, $\lambda_i^{(\kappa)}$, $\lambda_i^{(\epsilon)}$, and $\lambda_i^{(\delta)}$ for the coverage, adaptation, and acceptance constraints at the support point $\hat{x}_i$, and the Lagrangian subtracts the objective, which is to be maximized, and adds every constraint function, written such that it is positive exactly when the constraint is violated, weighted by its multiplier:

$$\begin{aligned}\mathcal{L}(\theta, \lambda) \;=\; -\mathrm{ESJD}(\theta) \;+\; \sum_{i=1}^{N} \Big[\; \lambda_i^{(\kappa)} \Big(\mathrm{KL}\left(\hat{p} \| q_\theta(\cdot|\hat{x}_i)\right) - \kappa \Big) \;+\; \lambda_i^{(\epsilon)} \Big(\mathrm{F}\left(q_\theta(\cdot|\hat{x}_i) \| p\right) - \epsilon \Big) \\ +\; \lambda_i^{(\delta)} \Big(\delta - \mathbb{E}_{v \sim p_v}\left[\alpha(\hat{x}_i, v)\right] \Big) \;\Big].\end{aligned} \tag{14}$$

Algorithm 1 alternates a stochastic gradient step on $\mathcal{L}$ over θ with the multipliers held fixed and a projected gradient ascent step over the multipliers, so that a multiplier grows while its constraint is

Algorithm 1 Primal-dual constrained learning of CI-MCMC

Inputs: unnormalized target $\tilde{p}$, thresholds $(\kappa, \epsilon, \delta)$, step sizes $\eta_\theta, \eta_\lambda > 0$, dual period k, iterations T, batch sizes V, B

$(\mathcal{S}, \hat{p}) \leftarrow \textsc{SupportSearch}(\tilde{p})$ (Algorithm 2, Appendix A.2)
Initialize: θ, $\lambda \leftarrow 0$, $\theta^\star \leftarrow \theta$, $J^\star \leftarrow -\infty$
for $t = 1, \dots, T$ **do**
  for $i = 1, \dots, N$ **do**
    draw $v_1, \dots, v_V \sim p_v$ and $y_1, \dots, y_B \sim \hat{p}$ ▷ batch
    $x'_s \leftarrow f_\theta^{(x)}(\hat{x}_i, v_s)$, $\alpha_s \leftarrow \alpha(\hat{x}_i, v_s)$ for $s = 1, \dots, V$ ▷ proposals
    $J_i \leftarrow \omega_i \frac{1}{V} \sum_{s=1}^{V} \|x'_s - \hat{x}_i\|^2 \alpha_s$ ▷ objective
    $s_i^\kappa \leftarrow \frac{1}{B} \sum_{b=1}^{B} \left[\log \hat{p}(y_b) - \log q_\theta(y_b \mid \hat{x}_i)\right] - \kappa$ ▷ coverage
    $s_i^\epsilon \leftarrow \frac{1}{V} \sum_{s=1}^{V} \left\|\nabla \log q_\theta(x'_s \mid \hat{x}_i) - \nabla \log \tilde{p}(x'_s)\right\|^2 - \epsilon$ ▷ adaptation
    $s_i^\delta \leftarrow \delta - \frac{1}{V} \sum_{s=1}^{V} \alpha_s$ ▷ acceptance
  end for
  $\mathcal{L} \leftarrow -\sum_i J_i + \sum_i \left(\lambda_i^\kappa s_i^\kappa + \lambda_i^\epsilon s_i^\epsilon + \lambda_i^\delta s_i^\delta\right)$ (Eq. (14))
  $\theta \leftarrow \theta - \eta_\theta \nabla_\theta \mathcal{L}$ ▷ primal step
  if $t \bmod k = 0$ **then**
    $\lambda_i^c \leftarrow \left[\lambda_i^c + \eta_\lambda \, s_i^c\right]_+$ for $c \in \{\kappa, \epsilon, \delta\}$, $i = 1, \dots, N$ ▷ dual step
  end if
  if $s_i^c \le 0$ for all c, i and $\sum_i J_i > J^\star$ **then**
    $\theta^\star \leftarrow \theta$, $J^\star \leftarrow \sum_i J_i$ ▷ best feasible iterate
  end if
end for
Output: $\theta^\star$

violated and returns to zero once the constraint holds; Chamon et al. (2023) show that for sufficiently expressive parametrizations a related iteration (where the primal step is replaced by a minimization oracle) yields a near-optimal and near-feasible solution of Eq. (11). All expectations are empirically estimated using a fresh batch per iteration with the smooth acceptance test of Section 3.3.

Inference. At test time CI-MCMC is an ordinary involutive MCMC chain: from the current state, it draws $v \sim p_v$, applies f_θ, and accepts according to the Metropolis–Hastings test in Eq. (4). Note that we do not use the smooth acceptance test from training during inference and discard the auxiliary variable of every accepted state. The cost per step is one forward pass of the network and one evaluation of the target density, the same as for any other involutive MCMC method.

5 Related Work

Classical samplers. Random walks (Metropolis et al., 1953) and discretized Langevin (Roberts & Tweedie, 1996b) and Hamiltonian (Duane et al., 1987; Homan & Gelman, 2014) dynamics are efficient on log-concave targets (Roberts & Tweedie, 1996a; Livingstone et al., 2019), but their proposals are local, so crossing the low-density region between two well-separated modes has exponentially small probability. Tempering (Swendsen & Wang, 1986; Geyer, 1991; Earl & Deem, 2005) reduces the separation by flattening the target at the cost of many auxiliary chains. JAMS (Pompe et al., 2020) instead locates the modes by optimization beforehand and constructs explicit jumps between them through affine maps. In benchmarks Łatuszyński et al. (2026), JAMS performed best.

Learned samplers. Adversarial objectives (Song et al., 2017; Spanbauer et al., 2020; Egorov et al., 2024) are popular to train neural parameterizations of MCMC kernels. Since ground-truth samples are inaccessible by definition, the *Markov GAN objective*(Song et al., 2017) bootstraps several chains to generate pseudo-samples from the target, rendering the overall objective brittle (Levy et al., 2018; Gabrié et al., 2022). A related use of discriminators is to estimate the acceptance ratio of an implicit proposal (Neklyudov et al., 2019; Turner et al., 2019). Non-adversarial objectives include the decorrelation of generated samples (Pasarica & Gelman, 2010; Levy et al., 2018), the maximization of proposal entropy (Hirt et al., 2021), reinforcement learning (Wang et al., 2025), and minimization

Table 1: Sliced Wasserstein-1 distance SW_1 (lower is better). Mean $\pm$ standard deviation over 100 repetitions. Best method per row in bold.

	d	CI-MCMC	JAMS	NeuTra	L2HMC	A-NICE	iMCMC	LRDS	LV-DDS
Gaussian	2	**0.55** $\pm$ 0.23	0.65 $\pm$ 0.24	17.88 $\pm$ 8.48	17.63 $\pm$ 8.25	5.49 $\pm$ 8.13	2.19 $\pm$ 1.65	1.23 $\pm$ 0.79	16.09 $\pm$ 5.79
	16	**0.45** $\pm$ 0.19	0.53 $\pm$ 0.19	15.32 $\pm$ 5.33	15.17 $\pm$ 5.16	10.89 $\pm$ 4.85	13.54 $\pm$ 6.28	4.77 $\pm$ 4.49	13.18 $\pm$ 5.63
	32	**0.57** $\pm$ 0.39	0.60 $\pm$ 0.23	15.53 $\pm$ 6.13	15.37 $\pm$ 6.04	15.60 $\pm$ 6.33	15.87 $\pm$ 6.06	4.44 $\pm$ 5.06	13.33 $\pm$ 7.17
	64	**0.70** $\pm$ 0.41	0.78 $\pm$ 0.29	15.47 $\pm$ 5.24	15.47 $\pm$ 5.15	16.03 $\pm$ 5.26	15.99 $\pm$ 5.20	4.36 $\pm$ 3.85	13.56 $\pm$ 4.36
Anharmonic	2	**0.19** $\pm$ 0.04	0.21 $\pm$ 0.04	4.65 $\pm$ 2.22	4.92 $\pm$ 1.77	0.83 $\pm$ 1.22	0.31 $\pm$ 0.13	0.32 $\pm$ 0.28	3.77 $\pm$ 2.35
	16	0.23 $\pm$ 0.05	**0.20** $\pm$ 0.03	4.47 $\pm$ 0.76	4.46 $\pm$ 0.75	2.01 $\pm$ 0.78	2.27 $\pm$ 0.88	0.71 $\pm$ 1.00	3.82 $\pm$ 1.18
	32	**0.22** $\pm$ 0.04	0.24 $\pm$ 0.04	4.48 $\pm$ 0.84	4.45 $\pm$ 0.86	3.46 $\pm$ 1.02	4.41 $\pm$ 1.27	1.25 $\pm$ 0.92	4.01 $\pm$ 1.01
	64	**0.31** $\pm$ 0.25	0.45 $\pm$ 0.22	4.45 $\pm$ 0.70	4.48 $\pm$ 0.70	4.88 $\pm$ 0.78	5.06 $\pm$ 0.71	2.00 $\pm$ 0.78	3.85 $\pm$ 0.96
Rosenbrock	2	2.79 $\pm$ 2.10	3.62 $\pm$ 1.77	4.76 $\pm$ 1.86	4.92 $\pm$ 1.63	0.31 $\pm$ 0.15	**0.28** $\pm$ 0.09	0.70 $\pm$ 1.05	4.17 $\pm$ 1.97
	16	**0.29** $\pm$ 0.11	0.61 $\pm$ 0.67	4.73 $\pm$ 1.02	4.66 $\pm$ 0.90	2.50 $\pm$ 1.28	3.05 $\pm$ 1.09	3.37 $\pm$ 2.27	3.95 $\pm$ 1.13
	32	**0.42** $\pm$ 0.06	0.68 $\pm$ 0.12	4.61 $\pm$ 0.95	4.65 $\pm$ 0.90	3.76 $\pm$ 1.12	4.38 $\pm$ 1.03	4.12 $\pm$ 1.36	4.09 $\pm$ 0.87
	64	**1.12** $\pm$ 0.26	1.37 $\pm$ 0.25	4.75 $\pm$ 0.77	4.77 $\pm$ 0.73	5.11 $\pm$ 0.89	5.03 $\pm$ 0.62	4.65 $\pm$ 1.39	4.20 $\pm$ 0.84

of (mode-seeking) reverse KL (Hoffman et al., 2019; Grenioux et al., 2023) and bootstrapped forward KL (Samsonov et al., 2022) of variational approximations . Diffusion models can be used for inference by learning the drift of an SDE from a base to the target with log-variance losses (Zhang & Chen, 2021; Vargas et al., 2022; Richter & Berner, 2023). Most recently it has been combined (Noble et al., 2024) with the construction of a reference distribution from mode search.

6 Experiments

We evaluate our method on variety of multidimensional ($d \in [2, 64]$) inference problems from mixture distributions (Section 6.1), over scientific problems (Sections 6.2 and 6.3), to conditioning in generative models (Section 6.4). When we have ground truth samples, we evaluate the quality using (i) (sliced) Wasserstein-1 distance W_1, (ii) Maximum Mean Discrepancy (MMD), (iii) Kolmogorov-Smirnov (KS) statistic. Full details on our setup can be found in Appendix B.

6.1 Multidimensional Mixture Models

First, we evaluate on different mixture families (Gaussian, Anharmonic, Rosenbrock) with five components each. These artificial problems allow to control the problem dimension ($d \in \{16, 32, 64\}$) and the separation of the modes (see Appendix B.3). We illustrate their low-dimensional ($d = 2$) versions in Appendix B.3.2 and show that all methods are able to sample except for those explicitly designed for unimodal distributions. Differences arise in higher dimensions ($d \geq 16$), where it becomes increasingly difficult to cover modes.

Table 1 reports the results, averaged over 100 repetitions per setting. CI-MCMC attains the smallest distance to the ground truth in all but two settings, and for $d \geq 16$ the distances of the adversarially trained methods are at least four times larger than those of CI-MCMC. This finding is consistent with the KS and MMD metrics in Appendix B.3. Here we also show that CI-MCMC recovers the mode weights reliably, while the other methods fail to mix and thus have large errors.

6.2 Statistical Mechanics

In statistical mechanics, the ϕ^4 theory (Krumhansl & Schrieffer, 1975) yields a continuous energy function describing systems from magnets to liquid-gas critical points, to study phase transitions and critical phenomena. We take the setup of Noble et al. (2024) and study the one-dimensional discretized (lattice) ϕ^4 model which yields a $d = 32$ double-well potential with two separated modes and Dirichlet boundary conditions $x_0 = x_{d+1} = 0$. The target distribution is then given by

$$\ln \pi_h(x) = -\beta \left(\frac{ad}{2} \sum_{i=1}^{d+1} (x_i - x_{i-1})^2 + \frac{1}{4ad} \sum_{i=1}^{d} (1 - x_i^2)^2 + \frac{h}{ad} \sum_{i=1}^{d} x_i \right), \tag{15}$$

where β is the inverse temperature, a the lattice spacing, and h an external field that controls the relative weight of the two modes. It is challenging to infer the mode weights $w_+ := p(x_{d/2} > 0)$ and $w_- := p(x_{d/2} < 0)$ correctly (Noble et al., 2024).

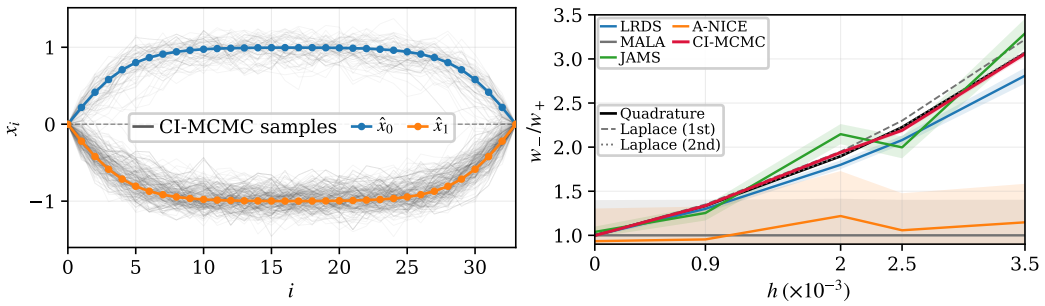

(a) Samples of a single CI-MCMC chain

(b) Estimated ratio of mode weights w_-/w_+

Figure 1: Lattice ϕ^4 model with samples obtained from CI-MCMC with $h = 0.035$ and $\mathcal{S} = \{\hat{x}_0, \hat{x}_1\}$ in 1a. A positive external field h tilts the double well towards the negative mode, so the mode-weight ratio w_-/w_+ in 1b grows with h. CI-MCMC estimate the ratio closest to quadrature.

In Fig. 1a we show samples from a *single* chain of CI-MCMC with $h = 0.035$ indicating that the chain has learned to mix between the modes. Figure 1b shows a sweep over h which changes the mode weights w_+ and w_-. Here, CI-MCMC yields samples that allow to estimate the mode-weight ratio w_-/w_+ closest to a reference computed by quadrature (see Appendix B.4 for details).

6.3 Reliability Engineering

Reliability engineering is about sampling failure states and estimating their probabilities to improve the reliability of complex systems. In particular, we study vehicle crash tests (Bichon et al., 2011) and sample $x \in \mathbb{R}^d$ from the space of $d = 11$ vehicle parameters (e.g. width of a beam) that caused the overall vehicle structure to fail under a side impact. Failure cases are rare and multimodal (see Appendix B.5), which renders standard MCMC methods ineffective. The *reliability sensitivity index (RSI)* (Xiao & Lu, 2017) describes how sensitive the reliability of the overall system is to changes in a system parameter x_i, as

$$\mathrm{RSI}_i = \frac{1}{2} \int \big| p(x_i) - p(x_i|\text{failure}) \big| \, dx_i \in [0, 1], \tag{16}$$

and Xiao & Nowak (2026) have computed Monte Carlo estimate from 100 runs of 10^7 samples each, which we consider the ground truth in Fig. 2 and shaded areas indicate the 99% quantile of their estimate. Inference methods sample $p(x_i|\text{failure})$, the distribution of the system parameter x_i conditioned on the overall system failing. Here, CI-MCMC yields the lowest error in estimating the RSI (6.5%) averaged over all system parameters.

6.4 Stanford Drone Dataset

Generative modeling estimates a data distribution $p(x)$ with NNs $q_\theta(x) \approx p(x)$ from samples allowing to generate new data points that are similar to the training data. However, conditioning such generative models on additional properties is hard. Special network architectures (Papamakarios et al., 2021; Vergari et al., 2021) have been proposed to allow efficient conditioning on *certain* properties. For *arbitrary* properties one typically resorts to MCMC methods. In particular, we predict trajectories $x \in \mathbb{R}^{60}$ of agents (pedestrians, bicyclists, cars, ...) in the Stanford university campus recorded by drones (Robicquet et al., 2016). We train a normalizing flow on $q_\theta(x) \approx p(x)$ on real trajectories (compare Fig. 3a). We pose a probabilistic inference problem by conditioning this generative model on (i) physical behavior by a bike model $p_{\text{phy}}(x)$ (compare Fig. 3b), and (ii) not violating the walls $\mathcal{W}$ annotated (Kurscheidt et al., 2025) in the dataset, and (iii) start (yellow circle) and end (dashed line around yellow star) positions of the trajectory. This yields the target distribution over trajectories x as

$$\pi(x) \propto q_\theta(x)\, p_{\text{phy}}(x)\, \mathbb{1}\{x \notin \mathcal{W}\}\; p_{\text{start}}(x)\, p_{\text{end}}(x). \tag{17}$$

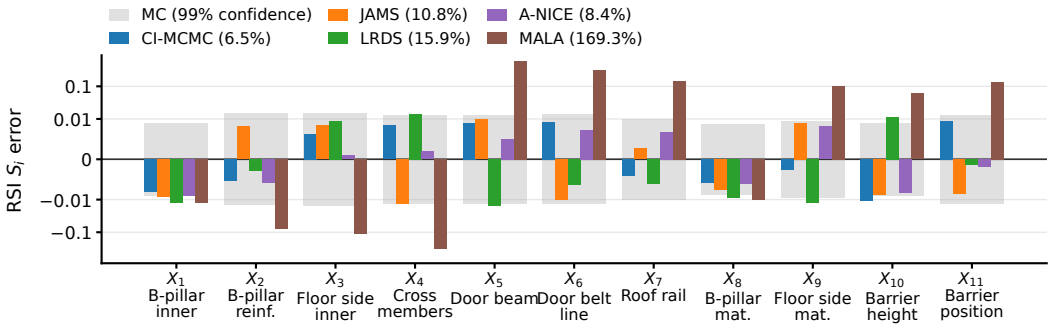

Figure 2: Reliability sensitivity index (RSI) for all system parameters x_i estimated from samples drawn from the failure distribution by different inference methods: Errors are estimated relative to the Monte Carlo estimate of Xiao & Nowak (2026). Shaded areas indicate the 99% quantile of their estimate. Error of methods averaged over all system parameters are shown in the legend.

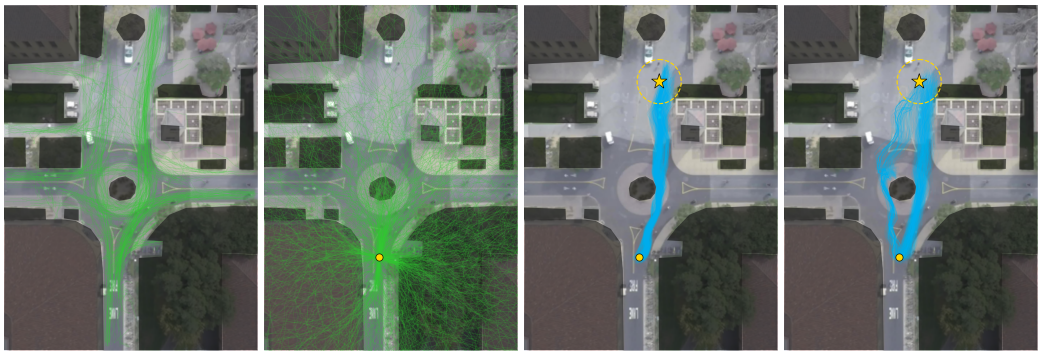

(a) Training data (b) Physics prior p_{phy} (c) JAMS (d) CI-MCMC

Figure 3: Conditional trajectory generation in a scene from the Stanford campus viewed from above. Dark polygons are the walls $\mathcal{W}$ annotated by Kurscheidt et al. (2025). 3a shows the trajectories used to train the generative model q_θ. 3b shows samples from the physical prior p_{phy}. 3d shows samples from CI-MCMC obtained from conditioning the generative model on the physical prior, the walls, and start (yellow dot) and end positions (yellow star). 3c shows samples from JAMS for the same target, which only cover the east route around the roundabout, failing to mix between the modes.

Rejection sampling $q_\theta(x)$ is not feasible as we empirically found that we required $\approx 10^7$ trials to obtain a single sample from $\pi(x)$. Furthermore, the target is multimodal as there are two routes around the roundabout and there are several samples in the training data that go west of the roundabout (see Appendix B.6 for details) that cause the generative model to have non-zero probabilities for such behaviors. The symmetry of the physical prior (compare Fig. 3b) increases the weight of this second mode. CI-MCMC is able to sample from both modes of $\pi(x)$ (compare Fig. 3d). JAMS, which has access to the same support points and Laplace approximations, fails to mix between the modes (compare Fig. 3c) as the two routes are not affine images of each other.

7 Conclusion

We proposed CI-MCMC, which trains involutive MCMC kernels by constrained learning: sampling efficiency is maximized subject to constraints that enforce coverage of all modes, adaptation to the target within each mode, and acceptance of proposals. We proved that any feasible solution yields a uniformly ergodic Markov chain, with a convergence rate that is explicit in the three constraint thresholds. On mixture models, a lattice ϕ^4 system, a reliability engineering problem, and conditional generative modeling, the trained chains sample all modes where classical and previously learned samplers do not.

AI USE STATEMENT

In this work, we used generative AI tools for editing and polishing the text, for checking proofs and references, and as a coding assistant. We have not used generative AI tools to generate research ideas or experimental results. We have reviewed all AI-assisted work: every proof step was verified by the authors, and all AI-assisted code was tested for correctness. We take responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

All assumptions and a complete proof of Theorem 1 are given in Appendix C. The training procedure is specified in Algorithm 1, the network architecture and the smooth acceptance test in Appendix A.1, the support point search in Appendix A.2, and the metrics and experimental setups in Appendix B.

REFERENCES

Christophe Andrieu, Nando De Freitas, Arnaud Doucet, and Michael I. Jordan. An introduction to MCMC for machine learning. *Machine learning*, 50(1):5–43, 2003. URL https://link.springer.com/article/10.1023/A:1020281327116.

D. Bakry and M. Émery. Diffusions hypercontractives. In Jacques Azéma and Marc Yor (eds.), *Séminaire de Probabilités XIX 1983/84*, volume 1123, pp. 177–206. Springer Berlin Heidelberg, Berlin, Heidelberg, 1985. ISBN 978-3-540-15230-9 978-3-540-39397-9. doi: 10.1007/BFb0075847. URL http://link.springer.com/10.1007/BFb0075847. Series Title: Lecture Notes in Mathematics.

Dominique Bakry, Ivan Gentil, and Michel Ledoux. *Analysis and geometry of Markov diffusion operators*, volume 348. Springer Science & Business Media, 2013. URL https://books.google.de/books?hl=de&lr=&id=gU3ABAAAQBAJ&oi=fnd&pg=PR4&dq=Analysis+and+geometry+of+Markov+diffusion+operators&ots=V0mS4GnHXs&sig=_k3uuE-jUWnM7PfOoNoNzQ7RzCg.

Barron J. Bichon, John M. McFarland, and Sankaran Mahadevan. Efficient surrogate models for reliability analysis of systems with multiple failure modes. *Reliability Engineering & System Safety*, 96(10):1386–1395, October 2011. ISSN 0951-8320. doi: 10.1016/j.ress.2011.05.008. URL https://www.sciencedirect.com/science/article/pii/S0951832011001062.

Chris Cannella and Vahid Tarokh. Semi-Empirical Objective Functions for MCMC Proposal Optimization. In *2022 26th International Conference on Pattern Recognition (ICPR)*, pp. 4758–4764, August 2022. doi: 10.1109/ICPR56361.2022.9956603. URL https://ieeexplore.ieee.org/document/9956603. ISSN: 2831-7475.

John Canny, Daniel Seita, and Anoop Korattikara. Fast Metropolis-Hastings and Natural Gradient. 2023. URL http://www.gatsby.ucl.ac.uk/~balaji/udl-camera-ready/UDL-23.pdf.

Luiz F. O. Chamon and Alejandro Ribeiro. Probably Approximately Correct Constrained Learning. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/c291b01517f3e6797c774c306591cc32-Abstract.html.

Luiz F. O. Chamon, Santiago Paternain, Miguel Calvo-Fullana, and Alejandro Ribeiro. Constrained Learning With Non-Convex Losses. *IEEE Trans. Inf. Theory*, 69(3):1739–1760, 2023. doi: 10.1109/TIT.2022.3187948.

Imre Csiszár. On information-type measure of difference of probability distributions and indirect observations. *Studia Sci. Math. Hungar.*, 2:299–318, 1967. URL https://cir.nii.ac.jp/crid/1571417124017192704.

Yilun Du and Igor Mordatch. Implicit Generation and Modeling with Energy Based Models. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://papers.neurips.cc/paper_files/paper/2019/hash/378a063b8fdb1db941e34f4bde584c7d-Abstract.html.

Simon Duane, Anthony D. Kennedy, Brian J. Pendleton, and Duncan Roweth. Hybrid monte carlo. *Physics letters B*, 195(2):216–222, 1987. URL https://www.sciencedirect.com/science/article/pii/037026938791197X.

David J. Earl and Michael W. Deem. Parallel tempering: Theory, applications, and new perspectives. *Physical Chemistry Chemical Physics*, 7(23):3910–3916, January 2005. ISSN 1463-9076. doi: 10.1039/b509983h. URL https://doi.org/10.1039/b509983h.

Evgenii Egorov, Riccardo Valperga, and Stratis Gavves. Ai-sampler: Adversarial Learning of Markov kernels with involutive maps. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 12304–12317. PMLR, July 2024. URL https://proceedings.mlr.press/v235/egorov24a.html.

Stefan C. Endres, Carl Sandrock, and Walter W. Focke. A simplicial homology algorithm for Lipschitz optimisation. *Journal of Global Optimization*, 72(2):181–217, October 2018. ISSN 1573-2916. doi: 10.1007/s10898-018-0645-y. URL https://doi.org/10.1007/s10898-018-0645-y.

Sam Foreman, Xiao-Yong Jin, and James C. Osborn. Deep Learning Hamiltonian Monte Carlo, May 2021. URL http://arxiv.org/abs/2105.03418. arXiv:2105.03418 [hep-lat].

Marylou Gabrié, Grant M. Rotskoff, and Eric Vanden-Eijnden. Adaptive Monte Carlo augmented with normalizing flows. *Proceedings of the National Academy of Sciences*, 119(10):e2109420119, March 2022. doi: 10.1073/pnas.2109420119. URL https://www.pnas.org/doi/abs/10.1073/pnas.2109420119.

Stuart Geman and Donald Geman. Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6):721–741, November 1984. ISSN 1939-3539. doi: 10.1109/TPAMI.1984.4767596. URL https://ieeexplore.ieee.org/document/4767596.

Charles J. Geyer. Markov Chain Monte Carlo Maximum Likelihood. 1991. URL https://hdl.handle.net/11299/58440.

Louis Grenioux, Alain Oliviero Durmus, Eric Moulines, and Marylou Gabrié. On Sampling with Approximate Transport Maps. In *Proceedings of the 40th International Conference on Machine Learning*, pp. 11698–11733. PMLR, July 2023. URL https://proceedings.mlr.press/v202/grenioux23a.html.

Marcel Hirt, Michalis Titsias, and Petros Dellaportas. Entropy-based adaptive hamiltonian monte carlo. *Advances in neural information processing systems*, 34:28482–28495, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/ef82567365bc6a0efb05d21837257424-Abstract.html.

Matthew Hoffman, Pavel Sountsov, Joshua V. Dillon, Ian Langmore, Dustin Tran, and Srinivas Vasudevan. NeuTra-lizing Bad Geometry in Hamiltonian Monte Carlo Using Neural Transport, March 2019. URL http://arxiv.org/abs/1903.03704. arXiv:1903.03704 [stat].

Matthew D. Hoffman and Pavel Sountsov. Tuning-Free Generalized Hamiltonian Monte Carlo. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, pp. 7799–7813. PMLR, May 2022. URL https://proceedings.mlr.press/v151/hoffman22a.html.

Matthew D. Homan and Andrew Gelman. The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *The Journal of Machine Learning Research*, 15(1):1593–1623, January 2014. ISSN 1532-4435. URL https://dl.acm.org/doi/10.5555/2627435.2638586.

Daphne Koller and Nir Friedman. *Probabilistic Graphical Models - Principles and Techniques*. MIT Press, 2009. ISBN 978-0-262-01319-2. URL http://mitpress.mit.edu/catalog/item/default.asp?ttype=2&tid=11886.

J. A. Krumhansl and J. R. Schrieffer. Dynamics and statistical mechanics of a one-dimensional model Hamiltonian for structural phase transitions. *Physical Review B*, 11(9):3535–3545, May 1975. doi: 10.1103/PhysRevB.11.3535. URL https://link.aps.org/doi/10.1103/PhysRevB.11.3535.

Leander Kurscheidt, Paolo Morettin, Roberto Sebastiani, Andrea Passerini, and Antonio Vergari. A Probabilistic Neuro-symbolic Layer for Algebraic Constraint Satisfaction. In *Proceedings of the Forty-first Conference on Uncertainty in Artificial Intelligence*, pp. 2431–2471. PMLR, July 2025. URL https://proceedings.mlr.press/v286/kurscheidt25a.html.

Daniel Levy, Matt D. Hoffman, and Jascha Sohl-Dickstein. Generalizing Hamiltonian Monte Carlo with Neural Networks. February 2018. URL https://openreview.net/forum?id=B1n8LexRZ.

Dong C. Liu and Jorge Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45(1):503–528, August 1989. ISSN 1436-4646. doi: 10.1007/BF01589116. URL https://doi.org/10.1007/BF01589116.

Samuel Livingstone, Michael Betancourt, Simon Byrne, and Mark Girolami. On the geometric ergodicity of Hamiltonian Monte Carlo. *Bernoulli*, 25(4A):3109–3138, November 2019. ISSN 1350-7265. doi: 10.3150/18-BEJ1083. URL https://projecteuclid.org/journals/bernoulli/volume-25/issue-4A/On-the-geometric-ergodicity-of-Hamiltonian-Monte-Carlo/10.3150/18-BEJ1083.full.

Richard McElreath. *Statistical rethinking: A Bayesian course with examples in R and Stan*. Chapman and Hall/CRC, 2018. URL https://www.taylorfrancis.com/books/mono/10.1201/9781315372495/statistical-rethinking-richard-mcelreath.

Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953. URL https://pubs.aip.org/aip/jcp/article-abstract/21/6/1087/202680.

Sean P. Meyn and Richard L. Tweedie. *Markov Chains and Stochastic Stability*. Communications and Control Engineering Series. Springer, 1993. ISBN 978-1-4471-3269-1. doi: 10.1007/978-1-4471-3267-7.

Kirill Neklyudov, Evgenii Egorov, and Dmitry P. Vetrov. The Implicit Metropolis-Hastings Algorithm. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 13932–13942, 2019. URL https://proceedings.neurips.cc/paper/2019/hash/80a160ff31266be2f93012a2a3eca713-Abstract.html.

Kirill Neklyudov, Max Welling, Evgenii Egorov, and Dmitry P. Vetrov. Involutive MCMC: a Unifying Framework. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 7273–7282. PMLR, 2020. URL http://proceedings.mlr.press/v119/neklyudov20a.html.

Maxence Noble, Louis Grenioux, Marylou Gabrié, and Alain Oliviero Durmus. Learned Reference-based Diffusion Sampler for multi-modal distributions. October 2024. URL https://openreview.net/forum?id=fmJUYgmMbL.

W. Nowak, F. P. J. de Barros, and Y. Rubin. Bayesian geostatistical design: Task-driven optimal site investigation when the geostatistical model is uncertain. *Water Resources Research*, 46(3), 2010. ISSN 1944-7973. doi: 10.1029/2009WR008312.

URL https://onlinelibrary.wiley.com/doi/abs/10.1029/2009WR008312. _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2009WR008312.

George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing Flows for Probabilistic Modeling and Inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021. ISSN 1533-7928. URL http://jmlr.org/papers/v22/19-1028.html.

Cristian Pasarica and Andrew Gelman. Adaptively scaling the Metropolis algorithm using expected squared jumped distance. *Statistica Sinica*, pp. 343–364, 2010. URL https://www.jstor.org/stable/24308995?casa_token=EWyMbqsNdzcAAAAA:BiWqT9o5IMKaaWBsFi1C0R-G7ustWrSPNL8q1yxGoEU6gxiUgoBhR-DvhlB2i3mWsuSn_gC3aUSXWra_MiYlYVjr69apcnH9kstJLWub5nCWz5U0Cuu3.

Emilia Pompe, Chris Holmes, and Krzysztof Łatuszyński. A framework for adaptive MCMC targeting multimodal distributions. *The Annals of Statistics*, 48(5):2930–2952, October 2020. ISSN 0090-5364, 2168-8966. doi: 10.1214/19-AOS1916. URL https://projecteuclid.org/journals/annals-of-statistics/volume-48/issue-5/A-framework-for-adaptive-MCMC-targeting-multimodal-distributions/10.1214/19-AOS1916.full.

Lorenz Richter and Julius Berner. Improved sampling via learned diffusions. October 2023. URL https://openreview.net/forum?id=h4pNROsO06.

Christian P. Robert and George Casella. *Monte Carlo Statistical Methods*. Springer Texts in Statistics. Springer New York, New York, NY, 2004. ISBN 978-1-4419-1939-7 978-1-4757-4145-2. doi: 10.1007/978-1-4757-4145-2. URL http://link.springer.com/10.1007/978-1-4757-4145-2.

G. O. Roberts and R. L. Tweedie. Geometric Convergence and Central Limit Theorems for Multidimensional Hastings and Metropolis Algorithms. *Biometrika*, 83(1):95–110, 1996a. ISSN 0006-3444. URL https://www.jstor.org/stable/2337435.

Gareth O. Roberts and Jeffrey S. Rosenthal. General state space Markov chains and MCMC algorithms. *Probability Surveys*, 1(none):20–71, January 2004. ISSN 1549-5787, 1549-5787. doi: 10.1214/154957804100000024. URL https://projecteuclid.org/journals/probability-surveys/volume-1/issue-none/General-state-space-Markov-chains-and-MCMC-algorithms/10.1214/154957804100000024.full.

Gareth O. Roberts and Richard L. Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341–363, December 1996b. ISSN 1350-7265. URL https://projecteuclid.org/journals/bernoulli/volume-2/issue-4/Exponential-convergence-of-Langevin-distributions-and-their-discrete-approximation bj/1178291835.full.

Alexandre Robicquet, Amir Sadeghian, Alexandre Alahi, and Silvio Savarese. Learning Social Etiquette: Human Trajectory Understanding In Crowded Scenes. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (eds.), *Computer Vision – ECCV 2016*, pp. 549–565, Cham, 2016. Springer International Publishing. ISBN 978-3-319-46484-8. doi: 10.1007/978-3-319-46484-8_33.

Sergey Samsonov, Evgeny Lagutin, Marylou Gabrié, Alain Durmus, Alexey Naumov, and Eric Moulines. Local-Global MCMC kernels: the best of both worlds. In *Advances in Neural Information Processing Systems*, volume 35, pp. 5178–5193. Curran Associates, Inc., 2022. doi: 10.52202/068431-0374. URL https://papers.nips.cc/paper_files/paper/2022/hash/21c86d5b10cdc28664ccdadf0a29065a-Abstract-Conference.html.

Tim Schneider, Amin Totounferoush, Wolfgang Nowak, and Steffen Staab. Probabilistic Regular Tree Priors for Scientific Symbolic Reasoning, June 2024. URL http://arxiv.org/abs/2306.08506. arXiv:2306.08506 [cs].

I. M Sobol'. On the distribution of points in a cube and the approximate evaluation of integrals. *USSR Computational Mathematics and Mathematical Physics*, 7(4):86–112, January 1967. ISSN 0041-5553. doi: 10.1016/0041-5553(67)90144-9. URL https://www.sciencedirect.com/science/article/pii/0041555367901449.

Jiaming Song, Shengjia Zhao, and Stefano Ermon. A-NICE-MC: Adversarial Training for MCMC. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 5140–5150, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/2417dc8af8570f274e6775d4d60496da-Abstract.html.

Span Spanbauer, Cameron E. Freer, and Vikash Mansinghka. Deep Involutive Generative Models for Neural MCMC. *CoRR*, abs/2006.15167, 2020. URL https://arxiv.org/abs/2006.15167. arXiv: 2006.15167.

R. H. Swendsen and J. S. Wang. Replica Monte Carlo simulation of spin glasses. *Physical Review Letters*, 57(21):2607–2609, November 1986. ISSN 1079-7114. doi: 10.1103/PhysRevLett.57.2607.

Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer series in statistics. Springer, 2009. ISBN 978-0-387-79051-0. doi: 10.1007/B13794. URL https://doi.org/10.1007/b13794.

Ryan Turner, Jane Hung, Eric Frank, Yunus Saatchi, and Jason Yosinski. Metropolis-Hastings Generative Adversarial Networks. In *Proceedings of the 36th International Conference on Machine Learning*, pp. 6345–6353. PMLR, May 2019. URL https://proceedings.mlr.press/v97/turner19a.html.

Francisco Vargas, Will Sussman Grathwohl, and Arnaud Doucet. Denoising Diffusion Samplers. September 2022. URL https://openreview.net/forum?id=8pvnfTAbu1f.

Santosh Vempala and Andre Wibisono. Rapid Convergence of the Unadjusted Langevin Algorithm: Isoperimetry Suffices. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper/2019/hash/65a99bb7a3115fdede20da98b08a370f-Abstract.html.

Antonio Vergari, YooJung Choi, Anji Liu, Stefano Teso, and Guy Van den Broeck. A Compositional Atlas of Tractable Circuit Operations for Probabilistic Inference. In *Advances in Neural Information Processing Systems*, volume 34, pp. 13189–13201. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper/2021/hash/6e01383fd96a17ae51cc3e15447e7533-Abstract.html.

Jasper A. Vrugt. Markov chain Monte Carlo simulation using the DREAM software package: Theory, concepts, and MATLAB implementation. *Environmental Modelling & Software*, 75:273–316, January 2016. ISSN 1364-8152. doi: 10.1016/j.envsoft.2015.08.013. URL https://www.sciencedirect.com/science/article/pii/S1364815215300396.

Congye Wang, Wilson Ye Chen, Heishiro Kanagawa, and Chris J. Oates. Reinforcement Learning for Adaptive MCMC. In *Proceedings of The 28th International Conference on Artificial Intelligence and Statistics*, pp. 640–648. PMLR, April 2025. URL https://proceedings.mlr.press/v258/wang25b.html.

Li K. Wenliang. Blindness of score-based methods to isolated components and mixing proportions. *CoRR*, abs/2008.10087, 2020. URL https://arxiv.org/abs/2008.10087. arXiv: 2008.10087.

Sinan Xiao and Zhenzhou Lu. Structural reliability sensitivity analysis based on classification of model output. *Aerospace Science and Technology*, 71:52–61, December 2017. ISSN 1270-9638. doi: 10.1016/j.ast.2017.09.009. URL https://www.sciencedirect.com/science/article/pii/S1270963817312713.

Sinan Xiao and Wolfgang Nowak. Reliability sensitivity analysis with multiple failure domains based on an extended two-stage Markov chain Monte Carlo simulation. *Reliability Engineering & System Safety*, 266:111718, February 2026. ISSN 0951-8320. doi: 10.1016/j.ress.2025.111718. URL https://www.sciencedirect.com/science/article/pii/S0951832025009184.

Qinsheng Zhang and Yongxin Chen. Path Integral Sampler: A Stochastic Control Approach For Sampling. October 2021. URL https://openreview.net/forum?id=_uCb2ynRu7Y.

Krzysztof Łatuszyński, Matthew T. Moores, and Timothée Stumpf-Fétizon. MCMC Methods for Multi-Modal Distributions. In *Handbook of Markov Chain Monte Carlo*. Chapman and Hall/CRC, 2 edition, 2026. Num Pages: 29.

A METHOD DETAILS

A.1 NETWORK ARCHITECTURE

As argued in the main text, the neural network architecture in CI-MCMC can not be arbitrary, but must satisfy the following properties: (i) it must be an involution, i.e. $f^{-1} = f$, (ii) it must be expressive enough to adapt to the target distribution within each mode, and (iii) it must allow for tractable evaluation of the marginal proposal density $q_\theta(x'|x)$. We realize these properties with the architecture depicted in Fig. 4. The overall network consists of three parts: $f_\theta = h \circ (g, g^{-1}) \circ h^{-1}$. This palindromic structure ensures that the overall network is an involution, since h^{-1} is the inverse of h and g^{-1} is the inverse of g and is inspired from Spanbauer et al. (2020). The inner bijection g is implemented as a stack of L coupling layers, each of which is an invertible transformation that splits the input into two halves, applies a monotone rational quadratic spline to one half conditioned on the other half, and then swaps the halves.

Tractable Marginal Densities. The final required property of the architecture is that it allows us to compute the marginal proposal density $q_\theta(x'|x)$ in closed form. Since the joint proposal is a point mass at $f_\theta(x, v)$, whose primary output is $x' = g^{-1}(h(x, v))$ with $h(x, v) = h_s(x) \odot v + h_t(x)$ (Fig. 4), the marginal follows from two changes of variables, $u = h(x, v)$ and $w = g^{-1}(u)$:

$$
\begin{aligned}
q_\theta(x'|x) &= \int \delta\big((x', v') - f_\theta(x, v)\big)\, p_v(v|x)\, dv\, dv' = \int \delta\big(x' - g^{-1}(h(x, v))\big)\, p_v(v|x)\, dv \\
&= \int \delta\big(x' - g^{-1}(u)\big)\, p_v\big((u - h_t(x)) \oslash h_s(x) \mid x\big) \left|\frac{\partial h(x, v)}{\partial v}\right|^{-1} du \\
&= \int \delta(x' - w)\, p_v\big((g(w) - h_t(x)) \oslash h_s(x) \mid x\big) \left|\frac{\partial g(w)}{\partial w}\right| \left|\frac{\partial h(x, v)}{\partial v}\right|^{-1} dw \\
&= p_v\big((g(x') - h_t(x)) \oslash h_s(x) \mid x\big) \left|\frac{\partial g(x')}{\partial x'}\right| \left|\frac{\partial h(x, v)}{\partial v}\right|^{-1},
\end{aligned}
\tag{18}
$$

where the first equality integrates out v', the second substitutes u with $dv = |\partial h/\partial v|^{-1} du$, the third substitutes w with $du = |\partial g/\partial w|\, dw$, and the last evaluates the remaining delta at $w = x'$. The Jacobian determinant of the affine coupling with respect to v is the product of the scales, $\prod_i h_s(x)_i$. This marginal proposal density can be efficiently evaluated by a forward pass through the neural networks h_s, h_t, g, computing the log-determinant of the Jacobian of g and evaluating the known distribution $p_v(v|x)$ of the auxiliary variable.

Smooth Acceptance Test. The acceptance probability of Eq. (4) is $\min\{1, r\}$ with

$$
r(x, v) = \frac{p(f(x, v))}{p(x, v)} \left|\frac{\partial f}{\partial (x, v)}\right|, \tag{19}
$$

whose gradient with respect to θ vanishes wherever $r > 1$. Following Canny et al. (2023), we replace it during training by the smooth acceptance test

$$
\alpha_k(x, v) = \big(1 + r(x, v)^{-k}\big)^{-1/k}, \tag{20}
$$

where $k > 0$ is a sharpness parameter. The test is differentiable everywhere, satisfies detailed balance for every k, recovers the Barker test $r/(1 + r)$ for $k = 1$, and approaches the Metropolis–Hastings test $\min\{1, r\}$ as $k \to \infty$. We use $k = 3$ in all experiments for the objective and the constraints of Algorithm 1, and switch back to the exact test $\min\{1, r\}$ at inference.

A.2 SUPPORT POINT SEARCH

Part of the training process is the support point search algorithm that is used to find the modes of the target distribution and construct the Laplace mixture reference distribution of Eq. (9). The search combines several steps, which are illustrated in Fig. 5. The combination of the steps yield Algorithm 2, the sub-routine called in the first line of Algorithm 1: We draw an initial set of candidate points from a Sobol sequence (Sobol', 1967), which covers the state space more evenly than uniform sampling (Fig. 5a). Following the *simplicial homology global optimization (SHGO)* algorithm

(Endres et al., 2018), we construct the candidate graph by connecting each point to its neighbors in space (Fig. 5b). The points whose log-density exceeds that of all their neighbors in the graph are considered candidates (Fig. 5c). Each candidate is refined by L-BFGS (Liu & Nocedal, 1989) on $\log p$, and the converged, deduplicated maxima form the support points $\mathcal{S}$ (Fig. 5d), from which the reference $\hat{p}$ of Eq. (9) is built. Finally, the algorithm returns the support points together with the Laplace mixture of Eq. (9) that serves as the reference distribution.

Algorithm 2 SUPPORTSEARCH: modes and reference distribution

Inputs: unnormalized target $\tilde{p}$, number of candidates $n_c \in \mathbb{N}$

1: $C \leftarrow$ first n_c points of a Sobol sequence
2: $C \leftarrow \{x \in C : \log \tilde{p}(x) \geq \log \tilde{p}(x') \text{ for all neighbors } x'\}$
3: $\mathcal{S} \leftarrow \text{unique}\big(\{\text{L-BFGS}(\log \tilde{p}, x) : x \in C\}\big)$
4: $\Sigma_i = \big(-\nabla^2 \log \tilde{p}(\hat{x}_i)\big)^{-1},\ \omega_i \propto \tilde{p}(\hat{x}_i)\det(2\pi\Sigma_i)^{1/2}$ for $\hat{x}_i \in \mathcal{S}$ (Eq. (9))

Output: $\mathcal{S},\ \hat{p} = \sum_i \omega_i \mathcal{N}(\cdot\,; \hat{x}_i, \Sigma_i)$

B EXPERIMENT DETAILS

B.1 METRICS

Our metrics evaluate two aspect of a method: overall sample quality and mode coverage.

B.1.1 SAMPLE QUALITY

When ground truth samples are available, we evaluate the quality of the samples drawn from a chain with several two-sample test metrics. Note that we report the metrics with respect to samples obtained from **single chains** of the respective methods to let the metrics reflect whether the method truely mixes. Let $\hat{\mu} = \frac{1}{N}\sum_{i=1}^{N} \delta_{x_i}$ denote the empirical distribution of N chain samples and $\hat{\nu} = \frac{1}{N}\sum_{j=1}^{N} \delta_{x_j^\star}$ that of N exact i.i.d. draws from the target.

Sliced Wasserstein-1 distance between two distributions α, β on $\mathbb{R}^d$ is

$$W_1(\alpha, \beta) = \inf_{\gamma \in \Gamma(\alpha,\beta)} \int \|x - x'\|_2 \, d\gamma(x, x'), \tag{21}$$

where $\Gamma(\alpha, \beta)$ is the set of couplings with marginals α and β. On the real line it has the closed form $W_1(\alpha, \beta) = \int_0^1 |F_\alpha^{-1}(q) - F_\beta^{-1}(q)|\, dq$. To keep this tractable in high dimensions we report the sliced Wasserstein-1 distance

$$\mathrm{SW}_1(\hat{\mu}, \hat{\nu}) = \mathbb{E}_{u \sim \mathcal{U}(\mathbb{S}^{d-1})}\Big[W_1\big(u_{\#}\hat{\mu},\, u_{\#}\hat{\nu}\big)\Big], \tag{22}$$

where $u_{\#}\hat{\mu}$ is the distribution of the projections $\langle u, x_i\rangle$ onto a direction u drawn uniformly from the unit sphere. The expectation is estimated with $L = 256$ random directions and the one-dimensional integral on a grid of 256 equally spaced quantile levels.

Maximum mean discrepancy for a positive definite kernel k is defined as

$$\mathrm{MMD}^2(\alpha, \beta) = \mathbb{E}\big[k(x, x')\big] - 2\,\mathbb{E}\big[k(x, x^\star)\big] + \mathbb{E}\big[k(x^\star, x^{\star\prime})\big], \tag{23}$$

with $x, x' \sim \alpha$ and $x^\star, x^{\star\prime} \sim \beta$ independent. We report the unbiased estimate

$$\widehat{\mathrm{MMD}}^2(\hat{\mu}, \hat{\nu}) = \frac{1}{N(N-1)}\sum_{i \neq i'} k(x_i, x_{i'}) - \frac{2}{N^2}\sum_{i,j} k(x_i, x_j^\star) + \frac{1}{N(N-1)}\sum_{j \neq j'} k(x_j^\star, x_{j'}^\star) \tag{24}$$

on 1500 samples per side, with the Gaussian kernel $k(x, x') = \exp\big(-\|x - x'\|_2^2 / 2h^2\big)$. The bandwidth h is chosen as the median pairwise distance of the *reference* sample $\hat{\nu}$.

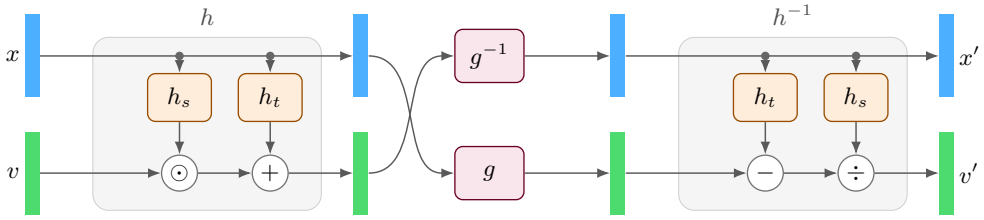

(a) Palindrome involution f_θ

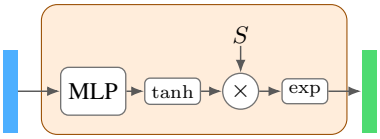

(b) Scale network h_s

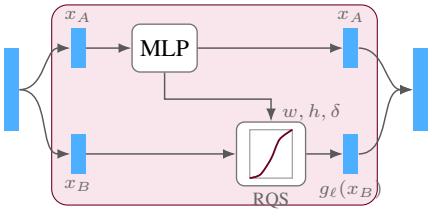

(c) One coupling layer of the inner flow g

Figure 4: *Palindrome involution.* (a) The affine coupling h scales v by $h_s(x)$ and shifts by $h_t(x)$, making the inner flow's input depend on x. The inner swap applies g^{-1} to the modified auxiliary and g to the original primary. Finally h^{-1} reverses the affine coupling. The symmetric $h \to (g^{-1}, g) \to h^{-1}$ layout is the palindrome structure that guarantees $f_\theta \circ f_\theta = \mathrm{id}$. (b) The scale network h_s: a spectrally normalized multilayer perceptron (MLP) whose output is squashed by $\tanh$, scaled by S, and exponentiated, so that $\log h_s(x) \in [-S, S]$; h_t has the same structure with output $B \tanh(\cdot)$. (c) One of the eight coupling layers of g, whose halves alternate between layers: the passive half x_A passes unchanged and conditions, through a spectrally normalized MLP, the K widths w, heights h, and knot slopes δ of a monotone rational quadratic spline applied elementwise to the active half x_B. The slopes are clipped to $[c_g, L_g]$.

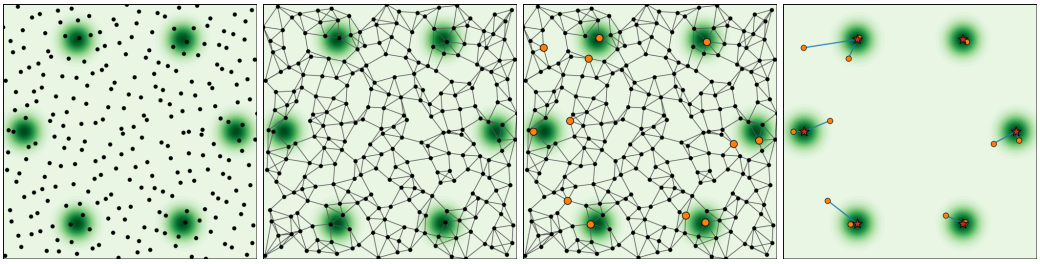

(a) Sobol points (b) Graph of neighbors (c) Candidates (d) L-BFGS

Figure 5: The support search algorithm involves four steps: (a) draw Sobol points, (b) construct a graph of neighbors, (c) keep candidates that are local maxima, and (d) refine candidates with L-BFGS to obtain the support points.

Kolmogorov–Smirnov statistic compares cumulative distribution functions in one dimension. We apply it to every coordinate marginal distribution and report the worst case,

$$\mathrm{KS}(\hat{\mu}, \hat{\nu}) = \max_{s=1,\dots,d} \sup_{t\in\mathbb{R}} \left| F^{(s)}_{\hat{\mu}}(t) - F^{(s)}_{\hat{\nu}}(t) \right|, \tag{25}$$

where $F^{(s)}_{\hat{\mu}}$ is the empirical cumulative distribution of the s-th coordinate of the chain samples.

B.1.2 Mode Coverage

Furthermore, we evaluate the ability of a method to reach all modes of a multimodal target distribution and estimate their relative weights. To do so, for a target with M modes at locations $\bar{x}_c$ with marginal standard deviations $\sigma_c \in \mathbb{R}^d$, we assign every sample to its nearest mode, as

$$c(x) = \operatorname*{arg\,min}_{c=1,\dots,M} \frac{\|x - \bar{x}_c\|_2}{\|\sigma_c\|_2}, \qquad \hat{w}_c = \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}\{c(x_i) = c\}\,. \tag{26}$$

The empirical mode weights $\hat{w}_c$ are compared to the true mode masses π_c by the total variation distance

$$\mathrm{TV}_{\text{mode}}(\hat{\mu}, p) = \frac{1}{2} \sum_{c=1}^{M} \left| \hat{w}_c - \pi_c \right|. \tag{27}$$

B.2 Baselines

MALA (Roberts & Tweedie, 1996b) proposes a local random step that follows the gradient of the target. JAMS (Pompe et al., 2020) locates the modes beforehand and jumps between them with fixed affine maps, mixing locally with MALA in between. A-NICE-MC (Song et al., 2017) trains a neural proposal adversarially against a discriminator that tells samples of the chain from bootstrapped samples of the target. iMCMC (Neklyudov et al., 2020) restricts the neural proposal to an involution and trains it with the same adversarial objective. L2HMC (Levy et al., 2018) learns a Hamiltonian-like dynamics by maximizing the expected jump distance of the chain. NeuTra (Hoffman et al., 2019) fits a normalizing flow to the target and runs HMC in the transformed space, where the geometry is simpler. LV-DDS (Vargas et al., 2022; Richter & Berner, 2023) learns the drift of a diffusion process that transports a Gaussian to the target and draws independent samples by simulating it. LRDS (Noble et al., 2024) replaces the Gaussian start of such a diffusion sampler by a mixture fitted around the modes so that all modes are reachable.

B.3 Mixture Models

Every target is a mixture $p(x) = \sum_{c=1}^{K} \omega_c\, p_c(x)$ of $K = 5$ components with weights ω_c and centers $\mu_c \in \mathbb{R}^d$, $d \in \{2, 16, 32, 64\}$. By exploiting the mixture structure with known modes weights, ground truth samples can be obtained for evaluations. To obtain diverse targets, we draw the centers μ_c uniformly from $[-3, 3]^d$ subject to a pairwise distance of at least 2, the weights are log-uniform with a ratio of at most 16 between the heaviest and the lightest component, and the variance σ^2 is log-uniform in $[0.05, 0.2]$ per target. The three families share the centers, the weights, and the variance in each coordinate, and differ only in the shape of the components. Two-dimensional versions are shown in Appendix B.3.2.

Gaussian. In the simplest case, the components are isotropic normal distributions,

$$p_c(x) \propto \exp\Big(- \frac{\|x - \mu_c\|_2^2}{2\sigma^2} \Big). \tag{28}$$

Anharmonic. To make the tails differ from a Gaussian, we let each component follow the Boltzmann law of an anharmonic oscillator, $\exp(-z^2/2 - \lambda z^4)$, adding no curvature at the mode but changing the shape away from it. It yields

$$p_c(x) \propto \exp\Big(- \sum_{i=1}^{d} \Big[\frac{(x_i - \mu_{c,i})^2}{2a_c^2} + \lambda_c \frac{(x_i - \mu_{c,i})^4}{a_c^4} \Big] \Big), \qquad a_c = \frac{\sigma}{\sqrt{m_2(\lambda_c)}}, \tag{29}$$

Table 2: Kolmogorov–Smirnov statistic KS (lower is better). Mean ± standard deviation over 100 repetitions. Best method per row in bold.

	d	CI-MCMC	JAMS	NeuTra	L2HMC	A-NICE	iMCMC	LRDS	LV-DDS
Gaussian	2	**0.030** ± 0.003	0.032 ± 0.003	0.660 ± 0.076	0.610 ± 0.064	0.166 ± 0.201	0.099 ± 0.034	0.058 ± 0.039	0.600 ± 0.107
	16	**0.029** ± 0.004	0.034 ± 0.003	0.656 ± 0.060	0.632 ± 0.043	0.492 ± 0.112	0.652 ± 0.129	0.199 ± 0.157	0.547 ± 0.137
	32	**0.033** ± 0.007	0.036 ± 0.002	0.626 ± 0.044	0.615 ± 0.037	0.738 ± 0.039	0.762 ± 0.021	0.173 ± 0.164	0.491 ± 0.138
	64	**0.041** ± 0.006	0.043 ± 0.002	0.637 ± 0.047	0.631 ± 0.053	0.771 ± 0.034	0.771 ± 0.034	0.172 ± 0.101	0.472 ± 0.083
Anharmonic	2	**0.031** ± 0.004	0.032 ± 0.002	0.486 ± 0.163	0.516 ± 0.088	0.089 ± 0.111	0.045 ± 0.010	0.038 ± 0.026	0.379 ± 0.187
	16	0.044 ± 0.009	**0.038** ± 0.003	0.549 ± 0.052	0.548 ± 0.052	0.281 ± 0.116	0.334 ± 0.115	0.094 ± 0.128	0.464 ± 0.126
	32	**0.041** ± 0.008	0.042 ± 0.003	0.534 ± 0.020	0.530 ± 0.019	0.491 ± 0.082	0.652 ± 0.095	0.151 ± 0.098	0.463 ± 0.105
	64	**0.049** ± 0.021	0.065 ± 0.015	0.538 ± 0.054	0.541 ± 0.051	0.736 ± 0.039	0.766 ± 0.031	0.245 ± 0.122	0.387 ± 0.077
Rosenbrock	2	0.265 ± 0.190	0.357 ± 0.148	0.500 ± 0.107	0.519 ± 0.066	0.043 ± 0.014	**0.042** ± 0.007	0.069 ± 0.086	0.438 ± 0.137
	16	**0.051** ± 0.016	0.084 ± 0.058	0.608 ± 0.076	0.583 ± 0.054	0.335 ± 0.143	0.427 ± 0.117	0.389 ± 0.222	0.486 ± 0.117
	32	**0.071** ± 0.010	0.097 ± 0.005	0.572 ± 0.039	0.580 ± 0.035	0.539 ± 0.097	0.650 ± 0.049	0.498 ± 0.136	0.469 ± 0.087
	64	0.184 ± 0.059	**0.179** ± 0.008	0.601 ± 0.036	0.606 ± 0.037	0.734 ± 0.024	0.742 ± 0.043	0.551 ± 0.115	0.416 ± 0.060

where the scale a_c keeps the variance the same with

$$m_2(\lambda) = \int z^2 e^{-z^2/2-\lambda z^4}\,\mathrm{d}z \,/ \int e^{-z^2/2-\lambda z^4}\,\mathrm{d}z. \tag{30}$$

For $\lambda_c = 0$, $m_2(0) = 1$ and the distribution is Gaussian. We draw λ_c log-uniformly from $[0.3, 10]$ per component, which renders Laplace approximations wrong by a factor between 2 and 10.

Rosenbrock. Finally, we rotate and bend components around their first coordinate inspired from the Rosenbrock function. With $y = D_c^{-1} Q_c^\top (x - \mu_c)$ and $u_i = y_i - s_c\,(y_1^2 - m_2(\lambda_c))$ for $i \geq 2$,

$$p_c(x) \propto \exp\left(-\frac{y_1^2}{2} - \lambda_c y_1^4 - \sum_{i=2}^{d} \left[\frac{u_i^2}{2} + \lambda_c u_i^4\right]\right), \tag{31}$$

where Q_c is a random orthogonal matrix that rotates each component by a quarter turn in one plane, the shear $s_c = \pm 0.3\sqrt{m_2(\lambda_c)}$ has a random sign per component and D_c scales to match the component variance to the Gaussian case.

Training. We search for support points on the box $[-4, 4]^d$ with 4096 Sobol points, five neighbors, and at most 20 points. The network follows Appendix A.1 with two hidden layers of 64 units. We train for 2000 primal steps with a batch of 4096 samples, a primal learning rate of 3×10^{-3}, a dual learning rate of 2×10^{-2}, a dual update every 100 primal steps, and keep the best feasible iterate.

Sampling. Every method runs 32 chains of 4000 steps per instance and sampler seed; the first 500 steps are discarded. All metrics of Appendix B.1 are computed per chain, averaged over the chains and over 10 different targets with 10 repetitions each.

B.3.1 Further results and metrics

Tables 2 and 3 report the Kolmogorov–Smirnov statistic and the squared maximum mean discrepancy for the same runs as Table 1. Both metrics show the same picture as the sliced Wasserstein distance: CI-MCMC and JAMS yield almost perfect samples, while the other methods fail to mix and thus have large errors.

Table 4 isolates the mode weights. CI-MCMC recovers them to within a few percent on the Gaussian and anharmonic mixtures in all dimensions and attains the smallest error in ten of the twelve settings, with JAMS close behind. NeuTra, L2HMC, and LV-DDS sit at errors of 0.7 to 0.8, which is the mass of the modes a chain never visits when it stays in the one it started in, and A-NICE-MC and iMCMC reach the true weights only in $d = 2$.

Table 5 restricts the comparison to samples within a mode, where the ranking changes. L2HMC attains the smallest within-mode distances on the anharmonic mixtures, and LRDS on the high-dimensional Rosenbrock mixtures, since a sampler that never leaves a mode spends all of its steps refining it. CI-MCMC and JAMS pay for their jumps with slightly larger within-mode errors, so mode jumping is what the learned proposal contributes, not local accuracy.

Table 3: Squared maximum mean discrepancy MMD^2 (lower is better). Mean $\pm$ standard deviation over 100 repetitions. Best method per row in bold.

	d	CI-MCMC	JAMS	NeuTra	L2HMC	A-NICE	iMCMC	LRDS	LV-DDS
Gaussian	2	**0.000** $\pm$ 0.000	0.000 $\pm$ 0.000	0.401 $\pm$ 0.157	0.402 $\pm$ 0.158	0.079 $\pm$ 0.160	0.006 $\pm$ 0.004	0.005 $\pm$ 0.009	0.379 $\pm$ 0.207
	16	**0.000** $\pm$ 0.000	0.000 $\pm$ 0.000	0.394 $\pm$ 0.105	0.394 $\pm$ 0.118	0.217 $\pm$ 0.100	0.281 $\pm$ 0.101	0.065 $\pm$ 0.091	0.304 $\pm$ 0.156
	32	**0.000** $\pm$ 0.000	0.000 $\pm$ 0.000	0.350 $\pm$ 0.059	0.344 $\pm$ 0.053	0.331 $\pm$ 0.058	0.346 $\pm$ 0.056	0.048 $\pm$ 0.069	0.263 $\pm$ 0.124
	64	**0.001** $\pm$ 0.000	0.001 $\pm$ 0.000	0.347 $\pm$ 0.080	0.350 $\pm$ 0.082	0.353 $\pm$ 0.082	0.352 $\pm$ 0.084	0.034 $\pm$ 0.037	0.279 $\pm$ 0.115
Anharmonic	2	**0.000** $\pm$ 0.000	0.000 $\pm$ 0.000	0.301 $\pm$ 0.145	0.345 $\pm$ 0.107	0.039 $\pm$ 0.090	0.002 $\pm$ 0.003	0.002 $\pm$ 0.004	0.196 $\pm$ 0.163
	16	0.001 $\pm$ 0.000	**0.000** $\pm$ 0.000	0.384 $\pm$ 0.114	0.384 $\pm$ 0.125	0.110 $\pm$ 0.071	0.115 $\pm$ 0.079	0.031 $\pm$ 0.085	0.284 $\pm$ 0.179
	32	**0.001** $\pm$ 0.000	0.001 $\pm$ 0.000	0.337 $\pm$ 0.052	0.332 $\pm$ 0.041	0.195 $\pm$ 0.054	0.275 $\pm$ 0.057	0.036 $\pm$ 0.040	0.286 $\pm$ 0.122
	64	**0.002** $\pm$ 0.002	0.003 $\pm$ 0.002	0.325 $\pm$ 0.089	0.329 $\pm$ 0.086	0.329 $\pm$ 0.097	0.349 $\pm$ 0.094	0.087 $\pm$ 0.111	0.235 $\pm$ 0.113
Rosenbrock	2	0.103 $\pm$ 0.107	0.166 $\pm$ 0.131	0.318 $\pm$ 0.111	0.344 $\pm$ 0.090	0.002 $\pm$ 0.003	**0.001** $\pm$ 0.001	0.012 $\pm$ 0.029	0.233 $\pm$ 0.147
	16	**0.002** $\pm$ 0.001	0.008 $\pm$ 0.015	0.400 $\pm$ 0.135	0.401 $\pm$ 0.138	0.143 $\pm$ 0.113	0.198 $\pm$ 0.135	0.253 $\pm$ 0.262	0.299 $\pm$ 0.166
	32	**0.004** $\pm$ 0.002	0.009 $\pm$ 0.002	0.336 $\pm$ 0.066	0.343 $\pm$ 0.071	0.219 $\pm$ 0.069	0.272 $\pm$ 0.049	0.279 $\pm$ 0.120	0.301 $\pm$ 0.126
	64	**0.027** $\pm$ 0.019	0.035 $\pm$ 0.004	0.340 $\pm$ 0.083	0.349 $\pm$ 0.083	0.358 $\pm$ 0.080	0.350 $\pm$ 0.095	0.348 $\pm$ 0.128	0.270 $\pm$ 0.082

Table 4: Mode-weight error $\mathrm{TV}_{\mathrm{mode}}$ (lower is better). Mean $\pm$ standard deviation over 100 repetitions. Best method per row in bold.

	d	CI-MCMC	JAMS	NeuTra	L2HMC	A-NICE	iMCMC	LRDS	LV-DDS
Gaussian	2	**0.019** $\pm$ 0.001	0.025 $\pm$ 0.001	0.797 $\pm$ 0.073	0.799 $\pm$ 0.075	0.188 $\pm$ 0.238	0.080 $\pm$ 0.033	0.057 $\pm$ 0.051	0.817 $\pm$ 0.116
	16	**0.020** $\pm$ 0.003	0.025 $\pm$ 0.001	0.790 $\pm$ 0.067	0.787 $\pm$ 0.058	0.539 $\pm$ 0.124	0.658 $\pm$ 0.131	0.260 $\pm$ 0.231	0.748 $\pm$ 0.173
	32	**0.024** $\pm$ 0.008	0.027 $\pm$ 0.001	0.779 $\pm$ 0.066	0.772 $\pm$ 0.068	0.757 $\pm$ 0.061	0.771 $\pm$ 0.066	0.216 $\pm$ 0.212	0.708 $\pm$ 0.212
	64	**0.031** $\pm$ 0.007	0.037 $\pm$ 0.002	0.787 $\pm$ 0.066	0.793 $\pm$ 0.068	0.791 $\pm$ 0.065	0.792 $\pm$ 0.069	0.230 $\pm$ 0.144	0.777 $\pm$ 0.136
Anharmonic	2	**0.020** $\pm$ 0.002	0.024 $\pm$ 0.001	0.663 $\pm$ 0.216	0.723 $\pm$ 0.122	0.096 $\pm$ 0.138	0.037 $\pm$ 0.012	0.040 $\pm$ 0.032	0.534 $\pm$ 0.264
	16	0.032 $\pm$ 0.008	**0.026** $\pm$ 0.001	0.790 $\pm$ 0.078	0.790 $\pm$ 0.075	0.355 $\pm$ 0.167	0.366 $\pm$ 0.121	0.138 $\pm$ 0.194	0.725 $\pm$ 0.163
	32	0.031 $\pm$ 0.007	**0.031** $\pm$ 0.001	0.769 $\pm$ 0.058	0.763 $\pm$ 0.057	0.542 $\pm$ 0.072	0.669 $\pm$ 0.085	0.213 $\pm$ 0.139	0.735 $\pm$ 0.168
	64	**0.044** $\pm$ 0.031	0.067 $\pm$ 0.025	0.772 $\pm$ 0.077	0.779 $\pm$ 0.072	0.755 $\pm$ 0.079	0.779 $\pm$ 0.073	0.376 $\pm$ 0.160	0.722 $\pm$ 0.169
Rosenbrock	2	0.363 $\pm$ 0.261	0.508 $\pm$ 0.185	0.691 $\pm$ 0.149	0.732 $\pm$ 0.106	0.039 $\pm$ 0.016	**0.034** $\pm$ 0.008	0.080 $\pm$ 0.120	0.623 $\pm$ 0.205
	16	**0.042** $\pm$ 0.012	0.093 $\pm$ 0.092	0.797 $\pm$ 0.058	0.797 $\pm$ 0.055	0.405 $\pm$ 0.181	0.482 $\pm$ 0.123	0.540 $\pm$ 0.298	0.734 $\pm$ 0.167
	32	**0.062** $\pm$ 0.011	0.108 $\pm$ 0.009	0.775 $\pm$ 0.065	0.774 $\pm$ 0.062	0.574 $\pm$ 0.121	0.664 $\pm$ 0.066	0.689 $\pm$ 0.185	0.728 $\pm$ 0.125
	64	**0.177** $\pm$ 0.064	0.235 $\pm$ 0.016	0.795 $\pm$ 0.057	0.794 $\pm$ 0.056	0.754 $\pm$ 0.071	0.761 $\pm$ 0.077	0.776 $\pm$ 0.139	0.794 $\pm$ 0.118

Table 5: Within-mode sliced Wasserstein-1 distance SW_1 (lower is better). Mean $\pm$ standard deviation over 100 repetitions. Best method per row in bold.

	d	CI-MCMC	JAMS	NeuTra	L2HMC	A-NICE	iMCMC	LRDS	LV-DDS
Gaussian	2	**0.121** $\pm$ 0.012	0.122 $\pm$ 0.011	0.548 $\pm$ 0.506	0.181 $\pm$ 0.124	0.259 $\pm$ 0.174	0.261 $\pm$ 0.084	0.172 $\pm$ 0.044	0.334 $\pm$ 0.163
	16	**0.122** $\pm$ 0.010	0.130 $\pm$ 0.011	0.385 $\pm$ 0.314	0.245 $\pm$ 0.092	0.642 $\pm$ 0.230	0.993 $\pm$ 0.163	0.165 $\pm$ 0.046	1.133 $\pm$ 0.621
	32	**0.131** $\pm$ 0.014	0.135 $\pm$ 0.010	0.253 $\pm$ 0.226	0.216 $\pm$ 0.149	1.083 $\pm$ 0.085	1.127 $\pm$ 0.007	0.165 $\pm$ 0.050	2.312 $\pm$ 1.138
	64	0.148 $\pm$ 0.011	**0.147** $\pm$ 0.008	0.246 $\pm$ 0.222	0.168 $\pm$ 0.152	1.134 $\pm$ 0.002	1.135 $\pm$ 0.003	0.168 $\pm$ 0.046	8.365 $\pm$ 3.772
Anharmonic	2	0.120 $\pm$ 0.014	0.119 $\pm$ 0.011	0.215 $\pm$ 0.257	**0.080** $\pm$ 0.010	0.151 $\pm$ 0.031	0.144 $\pm$ 0.013	0.112 $\pm$ 0.011	0.089 $\pm$ 0.018
	16	0.152 $\pm$ 0.017	0.143 $\pm$ 0.011	0.097 $\pm$ 0.013	**0.078** $\pm$ 0.017	0.337 $\pm$ 0.113	0.567 $\pm$ 0.149	0.105 $\pm$ 0.019	0.266 $\pm$ 0.118
	32	0.149 $\pm$ 0.023	0.157 $\pm$ 0.010	0.101 $\pm$ 0.009	**0.074** $\pm$ 0.010	0.716 $\pm$ 0.180	0.988 $\pm$ 0.147	0.115 $\pm$ 0.011	0.464 $\pm$ 0.177
	64	0.146 $\pm$ 0.019	0.186 $\pm$ 0.009	0.115 $\pm$ 0.012	**0.082** $\pm$ 0.017	1.081 $\pm$ 0.090	1.135 $\pm$ 0.004	0.098 $\pm$ 0.024	2.156 $\pm$ 0.596
Rosenbrock	2	0.150 $\pm$ 0.069	0.102 $\pm$ 0.029	0.192 $\pm$ 0.147	**0.091** $\pm$ 0.018	0.147 $\pm$ 0.021	0.140 $\pm$ 0.013	0.116 $\pm$ 0.013	0.092 $\pm$ 0.024
	16	0.168 $\pm$ 0.025	0.203 $\pm$ 0.008	0.752 $\pm$ 0.690	0.655 $\pm$ 0.331	0.582 $\pm$ 0.320	0.964 $\pm$ 0.538	**0.127** $\pm$ 0.031	0.212 $\pm$ 0.068
	32	**0.206** $\pm$ 0.018	0.252 $\pm$ 0.006	0.792 $\pm$ 0.788	0.894 $\pm$ 0.816	1.118 $\pm$ 0.502	1.444 $\pm$ 0.571	0.316 $\pm$ 0.546	0.535 $\pm$ 0.200
	64	0.384 $\pm$ 0.088	0.321 $\pm$ 0.016	1.114 $\pm$ 0.688	1.062 $\pm$ 0.473	1.765 $\pm$ 0.551	1.731 $\pm$ 0.440	**0.181** $\pm$ 0.030	2.295 $\pm$ 0.661

B.3.2 SAMPLE VISUALIZATION

Figure 6 shows the samples produced by every method on the two-dimensional instances of the three mixture families (Gaussian, anharmonic, Rosenbrock). Each column is one target and each row is one method; the top row shows the target density alone, so that every panel below it can be read against the density it is meant to reproduce. All panels of a column share the same axis limits.

B.4 STATISTICAL MECHANICS

The ϕ^4 model is discretized version of a free energy model over a real scalar field. It is the continuous counterpart of the Ising model and a standard model for the study of phase transitions (Krumhansl & Schrieffer, 1975). The instance we use was introduced by Gabrié et al. (2022); Noble et al. (2024) with the Dirichlet boundary conditions $x_0 = x_{d+1} = 0$, a discretization with $d = 32$, $a = 0.1$, and $\beta = 20$ as a benchmark for multimodal sampling. We adopt their setting. At $h = 0$ the target is symmetric under $x \mapsto -x$ and has two modes, a positive and a negative arch that interpolate between the boundary values and the phases ± 1 (Fig. 1a). The two are separated by a barrier whose height grows with β, so a local sampler started in one arch essentially never crosses to the other. We locate the modes by L-BFGS on $\ln \pi_h$, and use them as the support points. Following Noble et al. (2024), the mode weights are $w_+ = p(x_{d/2} > 0)$ and $w_- = p(x_{d/2} < 0)$, the probability that the middle site lies in the positive or the negative phase. A positive h lowers the energy of the negative phase, so w_-/w_+ grows with h.

Ground truth. Noble et al. (2024) compares against Laplace approximations at the two modes $\hat{x}_0$ and $\hat{x}_1$ in the notation of Eq. (9),

$$\frac{w_-}{w_+} \approx \frac{\pi_h(\hat{x}_1)}{\pi_h(\hat{x}_0)} \qquad \text{and} \qquad \frac{w_-}{w_+} \approx \frac{\omega_1}{\omega_0}, \tag{32}$$

referred to as zeroth and second order. For the one-dimensional chain, however, the mode weights can be computed exactly: since the energy only couples neighboring sites, the marginal of the middle site factorizes into d one-dimensional integrals, which we evaluate by quadrature on a grid (Krumhansl & Schrieffer, 1975).

Training and sampling. The network follows Appendix A.1 with two hidden layers of 64 units. We train once at $h = 0$ for 5000 primal steps with a batch of 256 samples, a primal learning rate of 3×10^{-4}, a dual learning rate of 10^{-2}, a dual update every 100 primal steps, and keep the best feasible iterate. The same trained involution is used for all values of h: only the acceptance test changes with the target π_h, so the tilt of the modes is accounted for exactly by the Metropolis–Hastings correction. For every h we run 32 chains of 5000 steps, half of them started in each mode, discard the first 500 steps, and assign each sample to the nearest mode, which agrees with the sign of the middle site. The ratio w_-/w_+ is the ratio of the assignment counts.

B.5 RELIABILITY ENGINEERING

The *vehicle side impact (VSI) problem* (Bichon et al., 2011) is a commonly used benchmark for system reliability with multiple failure modes. Full finite element simulation of a side impact crash are too expensive to be evaluated inside a probabilistic inference method. So six quantities that affect the safety of the vehicle occupants are represented by a polynomial approximation fitted to a small number of crash simulations, which are then used as a surrogate for the full simulation.

The input $x \in \mathbb{R}^{11}$ to the model summarizes the thicknesses of structural parts, material properties, and deviations of the impact barrier from its nominal position. All components are independent and normally distributed a priori with the means and standard deviations in Table 6. The six responses are the abdomen load L, the pubic symphysis force F, the upper rib deflection D_u, the lower viscous

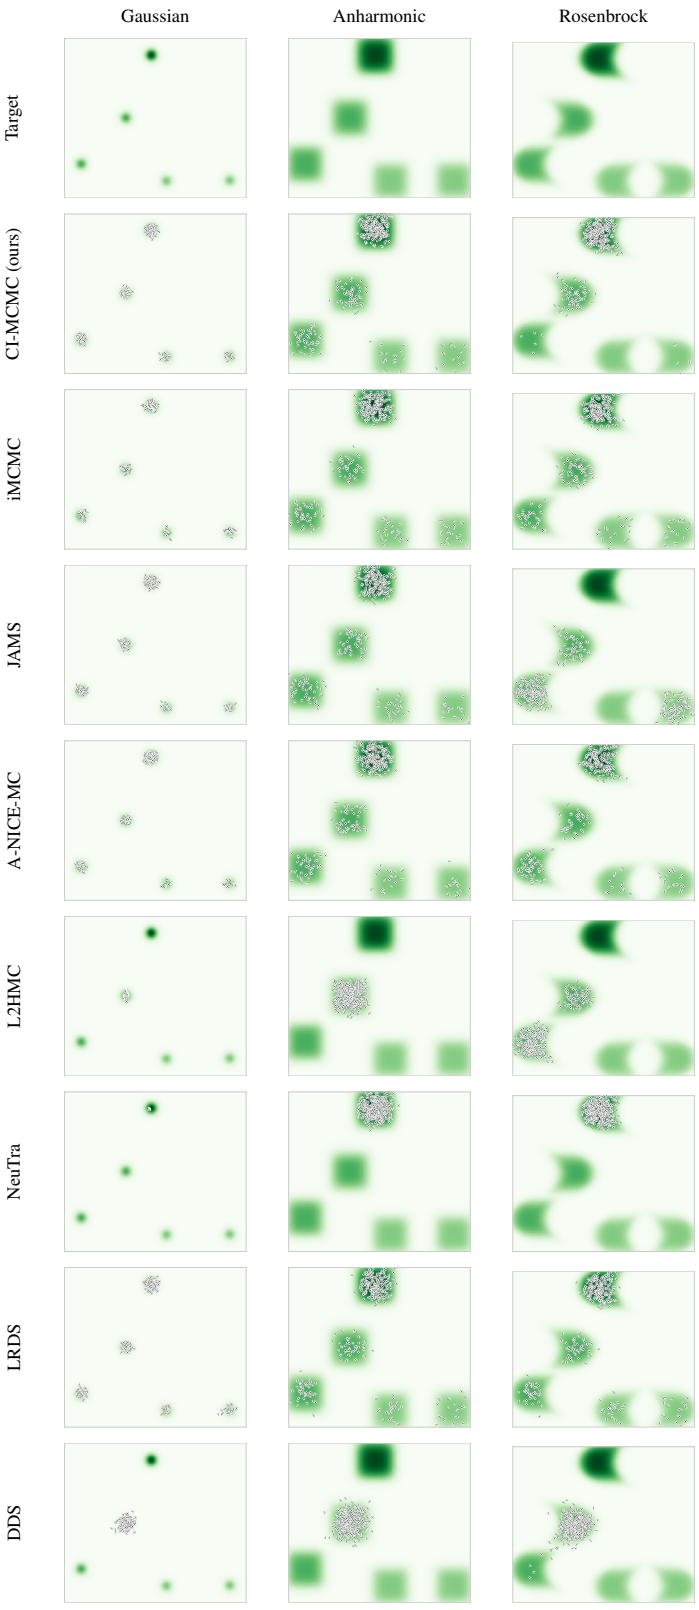

Figure 6: Samples on the two-dimensional mixture targets. Columns are targets, rows are methods; the top row shows the target density alone. Samples are drawn as black dots on top of the target density (green).

Table 6: Random variables of the vehicle side impact problem (Xiao & Nowak, 2026, Table 5).

Variable	Description	Mean	Std. dev.
x_1	B-pillar inner thickness	0.500	0.030
x_2	B-pillar reinforcement thickness	1.310	0.030
x_3	Floor side inner thickness	0.500	0.030
x_4	Cross members thickness	1.395	0.030
x_5	Door beam thickness	0.875	0.030
x_6	Door belt line reinforcement thickness	1.200	0.030
x_7	Roof rail thickness	0.400	0.030
x_8	B-pillar inner material	0.345	0.006
x_9	Floor side inner material	0.192	0.006
x_{10}	Barrier height	0.0	5.0
x_{11}	Barrier hitting position	0.0	5.0

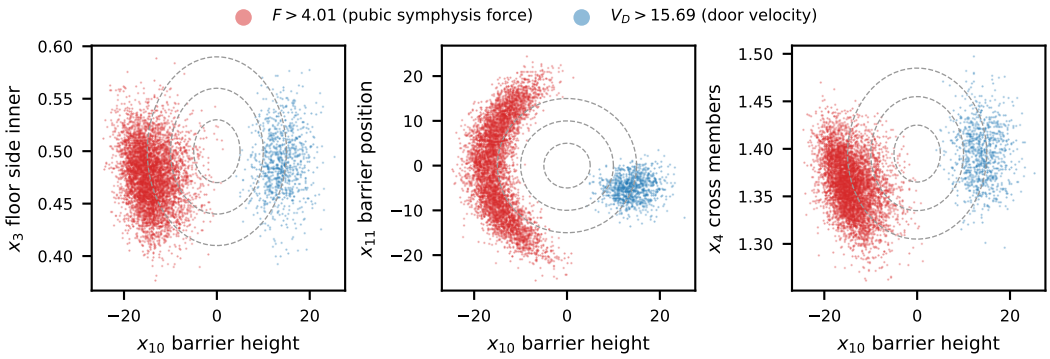

Figure 7: Failure samples colored by the violated criterion. Dashed ellipses mark one, two, and three standard deviations of the prior. The failure region consists of two disjoint domains in the barrier height x_{10}: a crescent at low barrier height where the pubic symphysis force exceeds its threshold, and a compact region at high barrier height where the door velocity does.

criterion VC_l, and the velocities of the B-pillar V_B and of the door V_D, given by

$$
\begin{aligned}
L &= 1.16 - 0.3717x_2x_4 - 0.00931x_2x_{10} - 0.484x_3x_9 + 0.01343x_6x_{10},\\
F &= 4.72 - 0.5x_4 - 0.19x_2x_3 - 0.0122x_4x_{10} + 0.009325x_6x_{10} + 0.000191x_{11}^2,\\
D_u &= 28.98 + 3.818x_3 - 4.2x_1x_2 + 0.0207x_5x_{10} + 6.63x_6x_9 - 7.7x_7x_8 + 0.32x_9x_{10},\\
VC_l &= 0.74 - 0.61x_2 - 0.163x_3x_8 + 0.001232x_3x_{10} - 0.166x_7x_9 + 0.227x_2^2,\\
V_B &= 10.58 - 0.674x_1x_2 - 1.95x_2x_8 + 0.02054x_3x_{10} - 0.0198x_4x_{10} + 0.028x_6x_{10},\\
V_D &= 16.45 - 0.489x_3x_7 - 0.843x_5x_6 + 0.0432x_9x_{10} - 0.0556x_9x_{11} - 0.000786x_{11}^2.
\end{aligned}
$$

The vehicle fails if *any* of these quantities exceeds the thresholds

$$L > 1.0 \ \vee \ F > 4.01 \ \vee \ D_u > 32.0 \ \vee \ VC_l > 0.32 \ \vee \ V_B > 9.9 \ \vee \ V_D > 15.69, \tag{33}$$

and the task is to sample from the distribution of the input parameters x conditioned on failure. The overall failure probability $P_F \approx 8.32 \times 10^{-4}$ (Xiao & Nowak, 2026) is small and the distribution is multimodal (see Fig. 7), which makes this a challenging problem for MCMC methods.

Training. The indicator of Eq. (33) has no gradient, so during training we replace it by the smooth surrogate $\sigma(-k\,G(x))$ with $k = 10$, which recovers the indicator as $k \to \infty$. At inference the exact indicator enters the acceptance test. The network follows Appendix A.1 with two hidden layers of 64 units. We train for 5000 primal steps with a batch of 256 samples, a primal learning rate of 3×10^{-4}, a dual learning rate of 10^{-2}, a dual update every 100 primal steps.

Reliability sensitivity. The reliability sensitivity index (Xiao & Lu, 2017) measures the influence of an input x_i on failure by the distance between its marginal prior and its marginal failure-conditional distribution. It is zero if failure does not change the distribution of x_i and one if failure

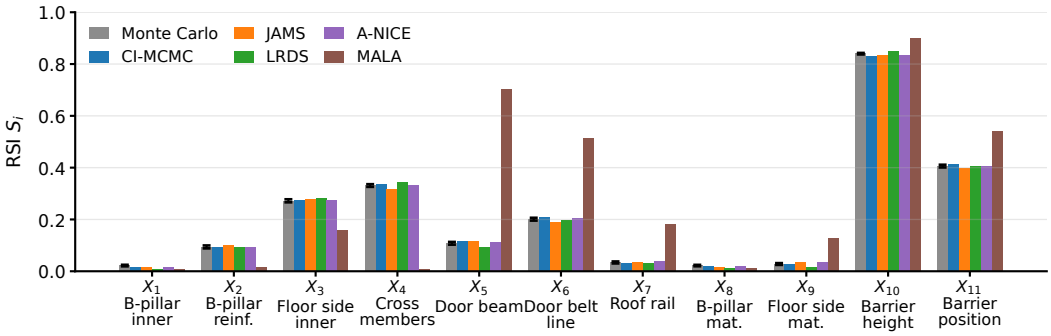

Figure 8: Reliability sensitivity index estimated by different methods.

fixes x_i deterministically. Figure 8 show that the methods tend to agree on the estimated sensitivity of the system parameters: Failure of the vehicle is by far most sensitive to the barrier height x_{10}, followed by the barrier hitting position x_{11}, the thickness of the cross members x_4, and the thickness of the floor side inner x_3. The material properties and the roof rail thickness are nearly irrelevant. Future work may use the CI-MCMC trained on the surrogate to estimate the RSI on expensive full finite element simulations, which would be infeasible with standard MCMC methods.

B.6 Stanford Drone Dataset

In this experiment we impose conditions on a *generative model* trained on the Stanford Drone Dataset (Robicquet et al., 2016). It contains drone videos of the Stanford campus annotated positions of pedestrians, cyclists, and vehicles. Kurscheidt et al. (2025) annotated the non-walkable area of two scenes as unions of convex polytopes. We use scene 12, a roundabout with a traffic island, 33 obstacle polytopes (Fig. 9d), and 290 training tracks (Fig. 9a).

Generative model. We consider trajectories of length $T = 30$ positions $x_t \in \mathbb{R}^2$ in pixel coordinates. To this data we fit a normalizing flow q_θ with rational quadratic spline coupling in six layers, trained by maximum likelihood. This is the generative model we use to impose further conditions.

Conditions. The first one is about physics. Since the data is sparse and noisy, we do not expect the flow to learn a physically plausible dynamics, so we add a physics prior p_{phy} that penalizes implausible accelerations and turning rates. Specifically, the physics prior is a kinematic bicycle model with acceleration a_t and turning rate ω_t obtained by finite differences of the positions, as

$$\log p_{\text{phy}}(x) = \sum_{t=1}^{T-2} \log \mathcal{N}\big((a_t, \omega_t); \mu, \Sigma\big) - \frac{1}{2s^2} \sum_{t=1}^{T-3} (\omega_{t+1} - \omega_t)^2, \tag{34}$$

with μ, Σ, and s estimated from the training tracks and the precision of ω scaled by 0.1. The second condition is about the environment: we require that the trajectory does not enter any of the 33 obstacle polytopes. The hard wall indicator is replaced by $\prod_{t,j} \sigma(-k \operatorname{pen}_j(x_t))$ with $k = 30$, where pen_j is a soft minimum of the halfspace residuals of obstacle j, evaluated at the positions. Draws from the flow alone (Fig. 9e) rarely satisfy walls, start, and end at once: rejection sampling accepted 2 of 1.1×10^7 draws.

Multimodality. The dataset contains few tracks that pass west of the traffic island (Fig. 9c), violating the traffic rules. The flow has learned this mode which renders the target multimodal, with two modes corresponding to the east and west routes. Here, our support search identifies $K = 6$ verified local maxima, three on each route. The network follows Appendix A.1 with two hidden layers of 64 units, trained for 4000 primal steps with a batch of 256, a primal learning rate of 3×10^{-3} with cosine decay, a dual learning rate of 2×10^{-2}, and the per-point acceptance level $\delta = 0.25$.

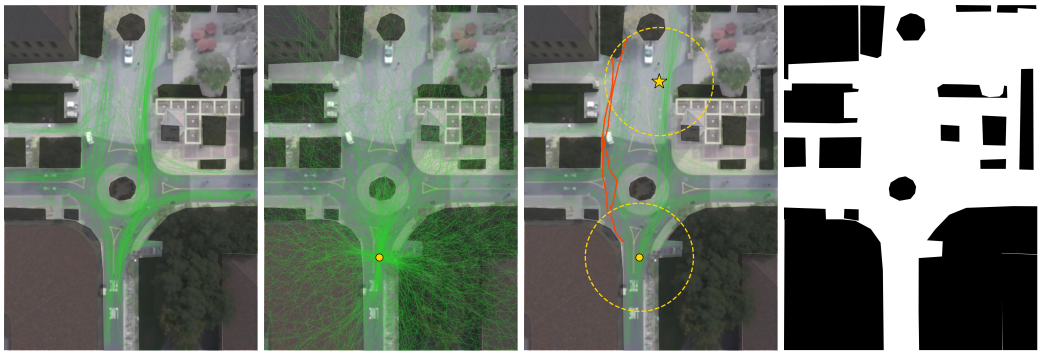

(a) Data (b) Physics prior (c) West route in the data (d) Walls



(e) Normalizing flow (f) MALA (g) JAMS (h) CI-MCMC

Figure 9: Conditioning a generative model on physics and environment constraints for a scene of the Stanford Drone Dataset. The flow alone (Fig. 9e) rarely satisfies all conditions (see walls Fig. 9d and physics Fig. 9b), MALA (Fig. 9f) never changes route, JAMS (Fig. 9g) has no jump accepted, and CI-MCMC (Fig. 9h) mixes between the two routes.

B.6.1 Results

Every panel of Fig. 9 shows samples from a single chain. CI-MCMC runs chains of $30{,}000$ steps at 4.7% acceptance. To improve mixing within the modes we alternate it with one MALA step. MALA alone (62% acceptance) never changes route. JAMS while having access to the same support points and Laplace approximations has none of its $12{,}000$ jump proposals accepted, because the routes are not affine images of each other, so its 30% acceptance is entirely local.

C Feasible Solutions Yield Ergodic Markov Chains

Throughout, Assumptions A1–A4 are in force and we use the notation of Section 3.4. On a set U_j covering a mode we denote by

$$\bar{q}_j(\cdot \mid x) := \frac{q_\theta(\cdot \mid x)}{\mathfrak{q}_\theta(U_j \mid x)}, \qquad \bar{p}_j := \frac{p}{\mathfrak{p}(U_j)} \tag{35}$$

the normalized restrictions of the two densities; both are well defined since $\mathfrak{q}_\theta(U_j \mid x) > 0$ by Assumption A1 and $\mathfrak{p}(U_j) > 0$ by positivity of p. Finally,

$$\epsilon_j(x) \;:=\; \int_{U_j} q_\theta(x' \mid x) \big\| \nabla \log q_\theta(x' \mid x) - \nabla \log p(x') \big\|^2 \, dx' \tag{36}$$

is the Fisher budget spent in U_j; since the cells are disjoint and the integrand is non-negative, $\sum_{j=1}^{M} \epsilon_j(x) \le \mathrm{F}\left(q_\theta(\cdot \mid x) \| p\right) \le \epsilon$. Recall that $\|\mu - \nu\|_{\mathrm{TV}} = \sup_B |\mu(B) - \nu(B)| = \frac{1}{2} \int |\mu - \nu|$ for distributions with densities (Scheffé's identity).

We use two classical inequalities, for whose proofs we refer to the literature.

Lemma 1 (Bakry–Émery criterion on a convex domain)**.** *Let $U \subset \mathbb{R}^d$ be open and convex and let μ be a probability density on U with $-\nabla^2 \log \mu \succeq mI$ on U for some $m > 0$. Then μ satisfies the logarithmic Sobolev inequality*

$$\mathrm{Ent}_\mu(f^2) \;\le\; \frac{2}{m}\,\mathbb{E}_\mu\|\nabla f\|^2 \qquad \textit{for all locally Lipschitz } f : U \to \mathbb{R}, \tag{37}$$

where $\mathrm{Ent}_\mu(g) := \mathbb{E}_\mu[g \log g] - \mathbb{E}_\mu[g] \log \mathbb{E}_\mu[g]$ denotes the entropy functional.

Proof. See Bakry & Émery (1985) and Bakry et al. (2013, Ch. 5). □

Corollary 1 (KL–Fisher inequality)**.** *Under the hypotheses of Lemma 1, for every positive C^1 probability density ν on U,*

$$\mathrm{KL}\,(\nu\|\mu) \;\le\; \frac{1}{2m}\,\mathrm{F}\,(\nu\|\mu)\,. \tag{38}$$

Proof. Step 1: choose the test function. Apply equation 37 to $f := \sqrt{\nu/\mu}$, which is positive and C^1 on U since both densities are (cf. Vempala & Wibisono, 2019, eq. (4)). Its square integrates to one under μ, $\mathbb{E}_\mu[f^2] = \int_U \nu = 1$, so the second term of the entropy vanishes and

$$\mathrm{Ent}_\mu(f^2) \;=\; \mathbb{E}_\mu\Big[\frac{\nu}{\mu}\log\frac{\nu}{\mu}\Big] \;=\; \int_U \nu \log\frac{\nu}{\mu} \;=\; \mathrm{KL}\,(\nu\|\mu)\,. \tag{39}$$

Step 2: the energy is a quarter of the Fisher divergence. By the chain rule, $\nabla f = \nabla\sqrt{\nu/\mu} = \frac{1}{2}\sqrt{\nu/\mu}\,\nabla\log(\nu/\mu)$, hence

$$\mathbb{E}_\mu\|\nabla f\|^2 \;=\; \frac{1}{4}\int_U \mu\,\frac{\nu}{\mu}\,\Big\|\nabla\log\frac{\nu}{\mu}\Big\|^2 \;=\; \frac{1}{4}\int_U \nu\,\big\|\nabla\log\nu - \nabla\log\mu\big\|^2 \;=\; \frac{1}{4}\,\mathrm{F}\,(\nu\|\mu)\,. \tag{40}$$

Step 3: combine. Inserting equation 39 and equation 40 into equation 37 gives equation 38. □

Lemma 2 (Pinsker's inequality)**.** *For any two probability distributions μ, ν with densities,*

$$\|\nu - \mu\|_{\mathrm{TV}} \;\le\; \sqrt{\tfrac{1}{2}\,\mathrm{KL}\,(\nu\|\mu)}. \tag{41}$$

Proof. See Csiszár (1967) or Tsybakov (2009, Lemma 2.5). □

Lemma 3 (Fisher divergence controls total variation on a cell)**.** *Let $U \subset \mathbb{R}^d$ be open and convex, let μ be a probability density on U with $-\nabla^2 \log \mu \succeq mI$ on U for some $m > 0$, and let ν be a positive C^1 probability density on U. Then*

$$\|\nu - \mu\|_{\mathrm{TV}} \;\le\; \sqrt{\mathrm{F}\,(\nu\|\mu)\,/(4m)}. \tag{42}$$

Proof. The density μ satisfies the hypotheses of Lemma 1, so Corollary 1 gives $\mathrm{KL}\,(\nu\|\mu) \le \mathrm{F}\,(\nu\|\mu)\,/(2m)$. Pinsker's inequality (Lemma 2) then yields

$$\|\nu - \mu\|_{\mathrm{TV}} \;\le\; \sqrt{\tfrac{1}{2}\,\mathrm{KL}\,(\nu\|\mu)} \;\le\; \sqrt{\frac{\mathrm{F}\,(\nu\|\mu)}{4m}}. \tag{43}$$

□

Lemma 4 (Coverage)**.** *For every $x \in \mathbb{R}^d$, $\sum_{c=0}^{M} \big|\mathfrak{q}_\theta(U_c \mid x) - \mathfrak{p}(U_c)\big| \le \tau + \sqrt{2\kappa}$.*

Proof. Step 1: from the KL constraint to total variation. Pinsker's inequality (Lemma 2) applied to the reference and the proposal, together with Assumption A4, gives

$$\big\|\widehat{\mathfrak{p}} - \mathfrak{q}_\theta(\cdot \mid x)\big\|_{\mathrm{TV}} \;\le\; \sqrt{\tfrac{1}{2}\,\mathrm{KL}\,(\hat{p}\|q_\theta(\cdot \mid x))} \;\le\; \sqrt{\kappa/2}. \tag{44}$$

Step 2: from total variation to the cells. For any two distributions μ, ν and the partition $U_0, \dots, U_M$, let B_+ be the union of the cells on which $\mu(U_c) > \nu(U_c)$ and B_- the union of the remaining cells. Then

$$\begin{aligned}\sum_{c=0}^{M} |\mu(U_c) - \nu(U_c)| \;&=\; \big(\mu(B_+) - \nu(B_+)\big) + \big(\nu(B_-) - \mu(B_-)\big) \\ &\le\; 2 \sup_{B} |\mu(B) - \nu(B)| \;=\; 2\, \|\mu - \nu\|_{\mathrm{TV}}\,,\end{aligned} \tag{45}$$

since the differences on the cells inside B_+ all have the same sign and add up to the difference on B_+, and likewise for B_-. With $\mu = \widehat{\mathfrak{p}}$ and $\nu = \mathfrak{q}_\theta(\cdot \mid x)$, equation 44 yields

$$\sum_{c=0}^{M} \big|\widehat{\mathfrak{p}}(U_c) - \mathfrak{q}_\theta(U_c \mid x)\big| \;\le\; \sqrt{2\kappa}. \tag{46}$$

This bounds the distance between the *reference* and the *proposal* on the cells.

Step 3: from the reference to the target. The bound equation 46 compares the proposal with the reference, whereas the claim compares the proposal with the target. Assumption A3 supplies the missing link, the distance between the reference and the target on the cells:

$$\sum_{c=0}^{M} \big|\widehat{\mathfrak{p}}(U_c) - \mathfrak{p}(U_c)\big| \;\le\; \tau. \tag{47}$$

For every cell c, adding and subtracting the reference mass $\widehat{\mathfrak{p}}(U_c)$ and applying the triangle inequality gives

$$\big|\mathfrak{q}_\theta(U_c \mid x) - \mathfrak{p}(U_c)\big| \;\le\; \big|\mathfrak{q}_\theta(U_c \mid x) - \widehat{\mathfrak{p}}(U_c)\big| \;+\; \big|\widehat{\mathfrak{p}}(U_c) - \mathfrak{p}(U_c)\big|. \tag{48}$$

Summing this inequality over $c = 0, \dots, M$ and bounding the two resulting sums by equation 46 and equation 47,

$$\sum_{c=0}^{M} \big|\mathfrak{q}_\theta(U_c \mid x) - \mathfrak{p}(U_c)\big| \;\le\; \sqrt{2\kappa} + \tau, \tag{49}$$

which is the claim. □

Lemma 5 (Proposal accuracy)**.** *For every $x \in \mathbb{R}^d$,*

$$2\, \|\mathfrak{q}_\theta(\cdot \mid x) - \mathfrak{p}\|_{\mathrm{TV}} \;\le\; \tau + \sqrt{2\kappa} + 2\eta + \sqrt{\epsilon/m}. \tag{50}$$

Proof. Step 1: total variation as an integral over cells. By Scheffé's identity, $2\, \|\mathfrak{q}_\theta(\cdot \mid x) - \mathfrak{p}\|_{\mathrm{TV}} = \int |q_\theta(\cdot \mid x) - p|$. The cells $U_0, \dots, U_M$ are disjoint and cover $\mathbb{R}^d$, so

$$2\, \|\mathfrak{q}_\theta(\cdot \mid x) - \mathfrak{p}\|_{\mathrm{TV}} \;=\; \sum_{c=0}^{M} \int_{U_c} |q_\theta(x' \mid x) - p(x')|\, dx'. \tag{51}$$

Step 2: the residual cell. Pointwise $|q_\theta(\cdot \mid x) - p| \le q_\theta(\cdot \mid x) + p$, hence $\int_{U_0} |q_\theta(\cdot \mid x) - p| \le \mathfrak{q}_\theta(U_0 \mid x) + \eta$, and since $\mathfrak{q}_\theta(U_0 \mid x) \le |\mathfrak{q}_\theta(U_0 \mid x) - \eta| + \eta$,

$$\int_{U_0} |q_\theta(x' \mid x) - p(x')|\, dx' \;\le\; \big|\mathfrak{q}_\theta(U_0 \mid x) - \mathfrak{p}(U_0)\big| + 2\eta. \tag{52}$$

Step 3: a set covering a mode, split into mass error and shape error. Fix $j \ge 1$. On U_j we have $q_\theta(\cdot \mid x) = \mathfrak{q}_\theta(U_j \mid x)\, \bar{q}_j(\cdot \mid x)$ and $p = \mathfrak{p}(U_j)\, \bar{p}_j$ by definition of the restrictions. Adding and subtracting $\mathfrak{q}_\theta(U_j \mid x)\, \bar{p}_j$,

$$q_\theta(\cdot \mid x) - p \;=\; \mathfrak{q}_\theta(U_j \mid x)\, (\bar{q}_j(\cdot \mid x) - \bar{p}_j) \;+\; \big(\mathfrak{q}_\theta(U_j \mid x) - \mathfrak{p}(U_j)\big)\, \bar{p}_j \qquad \text{on } U_j. \tag{53}$$

The triangle inequality under the integral, $\int_{U_j} \bar{p}_j = 1$, and Scheffé's identity for the restrictions give

$$\int_{U_j} |q_\theta(x' \mid x) - p(x')|\, dx' \;\le\; 2\, \mathfrak{q}_\theta(U_j \mid x)\; \|\bar{q}_j(\cdot \mid x) - \bar{p}_j\|_{\mathrm{TV}} \;+\; \big|\mathfrak{q}_\theta(U_j \mid x) - \mathfrak{p}(U_j)\big|. \tag{54}$$

Step 4: the shape error via Lemma 3. On U_j the restrictions differ from $q_\theta(\cdot \mid x)$ and p by constants, so $\nabla \log \bar{q}_j(\cdot \mid x) = \nabla \log q_\theta(\cdot \mid x)$ and $\nabla \log \bar{p}_j = \nabla \log p$, and therefore

$$\begin{aligned} \mathrm{F}\left(\bar{q}_j(\cdot \mid x) \| \bar{p}_j\right) &= \int_{U_j} \bar{q}_j(x' \mid x) \left\| \nabla \log q_\theta(x' \mid x) - \nabla \log p(x') \right\|^2 dx' \\ &= \frac{1}{\mathfrak{q}_\theta(U_j \mid x)} \int_{U_j} q_\theta(x' \mid x) \left\| \nabla \log q_\theta(x' \mid x) - \nabla \log p(x') \right\|^2 dx' \\ &= \frac{\epsilon_j(x)}{\mathfrak{q}_\theta(U_j \mid x)}. \end{aligned} \tag{55}$$

By Assumption A2, $\bar{p}_j$ is m-strongly log-concave on the open convex cell U_j, and $\bar{q}_j(\cdot \mid x)$ is positive and C^1 by Assumption A1, so Lemma 3 applies and yields $\|\bar{q}_j(\cdot \mid x) - \bar{p}_j\|_{\mathrm{TV}} \le \sqrt{\epsilon_j(x)/(4m\,\mathfrak{q}_\theta(U_j \mid x))}$. Multiplying by $2\,\mathfrak{q}_\theta(U_j \mid x)$, the proposal mass cancels the normalization of the restricted Fisher divergence:

$$2\,\mathfrak{q}_\theta(U_j \mid x)\, \|\bar{q}_j(\cdot \mid x) - \bar{p}_j\|_{\mathrm{TV}} \le 2\,\mathfrak{q}_\theta(U_j \mid x) \sqrt{\frac{\epsilon_j(x)}{4m\,\mathfrak{q}_\theta(U_j \mid x)}} = \sqrt{\frac{\mathfrak{q}_\theta(U_j \mid x)\,\epsilon_j(x)}{m}}. \tag{56}$$

Step 5: a set covering a mode, combined. Inserting equation 56 into equation 54 gives, for every $j \ge 1$,

$$\int_{U_j} |q_\theta(x' \mid x) - p(x')|\, dx' \le \left|\mathfrak{q}_\theta(U_j \mid x) - \mathfrak{p}(U_j)\right| + \sqrt{\frac{\mathfrak{q}_\theta(U_j \mid x)\,\epsilon_j(x)}{m}}. \tag{57}$$

Step 6: sum over the cells. Inserting equation 52 and equation 57 into equation 51,

$$2\, \|\mathfrak{q}_\theta(\cdot \mid x) - \mathfrak{p}\|_{\mathrm{TV}} \le \sum_{c=0}^{M} \left|\mathfrak{q}_\theta(U_c \mid x) - \mathfrak{p}(U_c)\right| + 2\eta + \sum_{j=1}^{M} \sqrt{\frac{\mathfrak{q}_\theta(U_j \mid x)\,\epsilon_j(x)}{m}}. \tag{58}$$

Lemma 4 bounds the first sum by $\tau + \sqrt{2\kappa}$.

Step 7: the shape sum via Cauchy–Schwarz. It remains to bound the last sum in equation 58. Writing each summand as a product of two square roots and applying the Cauchy–Schwarz inequality to the vectors $\left(\sqrt{\mathfrak{q}_\theta(U_j \mid x)}\right)_j$ and $\left(\sqrt{\epsilon_j(x)}\right)_j$,

$$\sum_{j=1}^{M} \sqrt{\mathfrak{q}_\theta(U_j \mid x)}\, \sqrt{\epsilon_j(x)} \le \Big(\sum_{j=1}^{M} \mathfrak{q}_\theta(U_j \mid x)\Big)^{1/2} \Big(\sum_{j=1}^{M} \epsilon_j(x)\Big)^{1/2}. \tag{59}$$

The first factor is at most one: the cells $U_1, \dots, U_M$ are disjoint, so $\sum_{j=1}^{M} \mathfrak{q}_\theta(U_j \mid x) = \mathfrak{q}_\theta\big(\bigcup_{j\ge1} U_j \,\big|\, x\big) \le 1$. The second factor is at most $\sqrt{\epsilon}$: by definition equation 36, $\sum_{j=1}^{M} \epsilon_j(x)$ is the integral of the non-negative function $q_\theta(x' \mid x)\|\nabla \log q_\theta(x' \mid x) - \nabla \log p(x')\|^2$ over the disjoint union $\bigcup_{j\ge1} U_j \subseteq \mathbb{R}^d$, hence at most its integral over $\mathbb{R}^d$, which is $\mathrm{F}\left(q_\theta(\cdot \mid x)\|p\right) \le \epsilon$ by Assumption A4. Dividing equation 59 by $\sqrt{m}$ therefore gives

$$\sum_{j=1}^{M} \sqrt{\frac{\mathfrak{q}_\theta(U_j \mid x)\,\epsilon_j(x)}{m}} \le \sqrt{\frac{\epsilon}{m}}, \tag{60}$$

and inserting this together with Lemma 4 into equation 58 yields the claim. □

Lemma 6 (One step)**.** *For every* $x \in \mathbb{R}^d$, $\|\mathfrak{t}_\theta(x, \cdot) - \mathfrak{p}\|_{\mathrm{TV}} \le \big(1 - \mathbb{E}_{v\sim p_v}\big[\alpha(x, v)\big]\big) + \|\mathfrak{q}_\theta(\cdot \mid x) - \mathfrak{p}\|_{\mathrm{TV}}$.

Proof. Step 1: split the proposal into its accepted and rejected part. Let $f_\theta^{(x)}$ denote the x-component of the involution, so that $\mathfrak{q}_\theta(\cdot \mid x)$ is the law of the proposed state $f_\theta^{(x)}(x, v)$ for $v \sim p_v$.

Weighting each proposed state by its acceptance probability or by its rejection probability defines the two measures $\mathfrak{a}$ and $\mathfrak{r}$,

$$\begin{aligned}\mathfrak{a}(x,B) &:= \int p_v(v)\,\alpha(x,v)\;\mathbb{1}\Big\{f_\theta^{(x)}(x,v)\in B\Big\}\;dv,\\ \mathfrak{r}(x,B) &:= \int p_v(v)\,\big(1-\alpha(x,v)\,\big)\,\mathbb{1}\Big\{f_\theta^{(x)}(x,v)\in B\Big\}\;dv,\end{aligned} \tag{61}$$

which add up to the proposal, $\mathfrak{q}_\theta(B \mid x) = \mathfrak{a}(x,B) + \mathfrak{r}(x,B)$ for every measurable B. Both are non-negative, and their total masses are the expected acceptance and rejection probabilities, $\mathfrak{a}(x,\mathbb{R}^d) = \mathbb{E}_{v\sim p_v}\big[\alpha(x,v)\big]$ and $\mathfrak{r}(x,\mathbb{R}^d) = 1-\mathbb{E}_{v\sim p_v}\big[\alpha(x,v)\big]$.

Step 2: write the kernel in the same terms. One step of the chain moves to the proposed state if it is accepted and stays at x otherwise, so

$$\mathfrak{t}_\theta(x,B) \;=\; \mathfrak{a}(x,B) + \big(1-\mathbb{E}_{v\sim p_v}\big[\alpha(x,v)\big]\big)\,\mathbb{1}\{x\in B\}\,. \tag{62}$$

The kernel and the proposal thus share the accepted part and differ only in where the rejected mass $1-\mathbb{E}_{v\sim p_v}\big[\alpha(x,v)\big]$ is placed: on the atom at x for the kernel, on the rejected proposals for the proposal.

Step 3: kernel and proposal are close. Subtracting the two decompositions, for every measurable B,

$$\big|\mathfrak{t}_\theta(x,B) - \mathfrak{q}_\theta(B\mid x)\big| \;=\; \big|\big(1-\mathbb{E}_{v\sim p_v}\big[\alpha(x,v)\big]\big)\,\mathbb{1}\{x\in B\} - \mathfrak{r}(x,B)\big| \;\le\; 1-\mathbb{E}_{v\sim p_v}\big[\alpha(x,v)\big], \tag{63}$$

because both terms inside the absolute value lie in $[0, 1-\mathbb{E}_{v\sim p_v}\big[\alpha(x,v)\big]]$. Taking the supremum over B gives $\|\mathfrak{t}_\theta(x,\cdot) - \mathfrak{q}_\theta(\cdot\mid x)\|_{\mathrm{TV}} \le 1-\mathbb{E}_{v\sim p_v}\big[\alpha(x,v)\big]$.

Step 4: triangle inequality. Total variation is a metric, so

$$\begin{aligned}\|\mathfrak{t}_\theta(x,\cdot) - \mathfrak{p}\|_{\mathrm{TV}} \;&\le\; \|\mathfrak{t}_\theta(x,\cdot) - \mathfrak{q}_\theta(\cdot\mid x)\|_{\mathrm{TV}} + \|\mathfrak{q}_\theta(\cdot\mid x) - \mathfrak{p}\|_{\mathrm{TV}}\\ &\le\; \big(1-\mathbb{E}_{v\sim p_v}\big[\alpha(x,v)\big]\big) + \|\mathfrak{q}_\theta(\cdot\mid x) - \mathfrak{p}\|_{\mathrm{TV}}\,.\end{aligned} \tag{64}$$

□

Proof of Theorem 1. Step 1: the target is invariant for the kernel. The involutive step leaves the joint distribution $p(x)\,p_v(v)$ invariant (Neklyudov et al., 2020), and resampling $v\sim p_v$ is a Gibbs step for the same joint distribution. One step of $\mathfrak{t}_\theta$ is the composition of the two, so if $x\sim\mathfrak{p}$ then the next state has law $\mathfrak{p}$ again: $\mathfrak{p}$ is invariant for $\mathfrak{t}_\theta$.

Step 2: the worst-case distance to the target is submultiplicative. Define the worst-case distance after n steps,

$$D(n) \;:=\; 2\,\sup_{x\in\mathbb{R}^d}\|\mathfrak{t}_\theta^n(x,\cdot) - \mathfrak{p}\|_{\mathrm{TV}}\,. \tag{65}$$

For a kernel with invariant distribution $\mathfrak{p}$ this quantity is submultiplicative (Roberts & Rosenthal, 2004, Prop. 3),

$$D(n_1+n_2) \;\le\; D(n_1)\,D(n_2) \qquad \text{for all } n_1,n_2\ge 1, \tag{66}$$

and in particular $D(n)\le D(1)^n$.

Step 3: the one-step distance is at most ρ. By Lemma 6, then Lemma 5, and finally Assumption A4 in the form $\mathbb{E}_{v\sim p_v}\big[\alpha(x,v)\big]\ge\delta$, for every x

$$\begin{aligned}2\;\|\mathfrak{t}_\theta(x,\cdot) - \mathfrak{p}\|_{\mathrm{TV}} \;&\le\; 2\Big(1-\mathbb{E}_{v\sim p_v}\big[\alpha(x,v)\big]\Big) + 2\;\|\mathfrak{q}_\theta(\cdot\mid x) - \mathfrak{p}\|_{\mathrm{TV}}\\ &\le\; 2\Big(1-\mathbb{E}_{v\sim p_v}\big[\alpha(x,v)\big]\Big) + \tau + \sqrt{2\kappa} + 2\eta + \sqrt{\epsilon/m}\\ &\le\; 2(1-\delta) + \tau + \sqrt{2\kappa} + 2\eta + \sqrt{\epsilon/m} \;=\; \rho.\end{aligned} \tag{67}$$

The right-hand side does not depend on x, so taking the supremum gives $D(1)\le\rho$.

Step 4: conclude. Combining the definition equation 65 with Steps 2 and 3, for every x and $n\ge 1$,

$$\|\mathfrak{t}_\theta^n(x,\cdot) - \mathfrak{p}\|_{\mathrm{TV}} \;\le\; \tfrac{1}{2}\,D(n) \;\le\; \tfrac{1}{2}\,D(1)^n \;\le\; \tfrac{1}{2}\,\rho^n, \tag{68}$$

which is the claim. □