Strictly Constrained Generative Modeling via Split Augmented Langevin Sampling

Matthieu Blanke Courant Institute New York University LEAP NSF STC mb10503@nyu.edu

Sara Shamekh Courant Institute New York University LEAP NSF STC ss18284@nyu.edu Yongquan Qu Columbia University LEAP NSF STC Pasteur Labs yq2340@columbia.edu

Pierre Gentine Columbia University LEAP NSF STC pg2328@columbia.edu

Abstract

Deep generative models hold great promise for representing complex physical systems, but their deployment is currently limited by the lack of guarantees on the physical plausibility of the generated outputs. Ensuring that known physical constraints are enforced is therefore critical when applying generative models to scientific and engineering problems. We address this limitation by developing a principled framework for sampling from a target distribution while rigorously satisfying physical constraints. Leveraging the variational formulation of Langevin dynamics, we propose Split Augmented Langevin (SAL), a novel primal-dual sampling algorithm that enforces constraints progressively through variable splitting, with convergence guarantees. While the method is developed theoretically for Langevin dynamics, we demonstrate its effective applicability to diffusion models. In particular, we use constrained diffusion models to generate physical fields satisfying energy and mass conservation laws. We apply our method to diffusion-based data assimilation on a complex physical system, where enforcing physical constraints substantially improves both forecast accuracy and the preservation of critical conserved quantities. We also demonstrate the potential of SAL for challenging feasibility problems in optimal control.

1 Introduction

Generative deep learning methods have recently emerged as powerful tools to model and sample from complex data distributions, with successful applications in image synthesis [1], protein and material design [2], and probabilistic weather forecasting [3]. By learning a stochastic process from a training dataset, these models can generate arbitrarily many plausible samples conditioned on partial information. They are particularly useful in the physical sciences, where data is often scarce and multiple states may be consistent with available observations [4, 5].

In domains such as image or text generation, visual or perceptual similarity is often sufficient to assess sample quality. Scientific applications, in contrast, frequently require samples to satisfy strict mathematical constraints, such as conservation laws or system dynamics [6]. In such cases, approximate resemblance is not enough: generated samples must obey the governing physical principles. This requirement becomes even more critical when generative models are used out of

distribution or in an autoregressive fashion, where small violations can accumulate and severely degrade long-term accuracy [7]. Developing constrained sampling methods applicable to pre-trained generative models in a zero-shot scenario (*i.e.* without additional training) is therefore crucial.

Modern generative models, including energy-based, score-based, and diffusion models [8–11], typically rely on gradually transforming noise into samples from the data distribution through a learned stochastic process. A core mechanism underlying many of these models is Langevin dynamics, where noisy gradient steps push the samples toward high-likelihood regions. Enforcing mathematical constraints during Langevin sampling remains a challenging problem. A natural idea is to project each iterate onto the constraint set, leading to the classical projected Langevin algorithm [12, 13] and, more recently, to projected diffusion [14]. Other approaches draw from constrained optimization, using barrier functions [15] or mirror maps [16]. While these methods offer theoretical guarantees in convex settings, they tend to perform poorly when applied to non-convex constraints, which are common in physical systems. In such cases, strict projections or barriers can cause the dynamics to become trapped in limited regions of the constraint set, hindering exploration and introducing significant sampling bias. In another line of work, Chamon et al. [17] propose a primal-dual Langevin algorithm, where constraints are satisfied in expectation with a Lagrangian relaxation. However, for many physical applications, enforcing strict constraints for all samples is essential. To the best of our knowledge, no existing method combines unbiased exploration with tight constraint satisfaction.

Contributions Inspired by variational formulations of Langevin dynamics and primal-dual optimization, we propose a novel sampling algorithm that bridges the gap between unbiased Langevin sampling and constraint satisfaction, called Split Augmented Langevin (SAL). Our method enforces hard constraints while preserving the exploration capability of Langevin dynamics. It guarantees strict constraint satisfaction and benefits from convergence guarantees via duality analysis. We show that our method generalizes to deep generative modeling and diffusion models. We demonstrate the effectiveness of SAL on physically-constrained sampling tasks on complex systems, including data assimilation problems where maintaining physical invariants is key to reliable forecasting, and on non-convex feasibility problems in optimal control.

2 Problem formulation of constrained Langevin sampling

In this section, we provide a mathematical formulation of constrained sampling: given a generative model and a constraint set, our goal is to generate samples from the conditional distribution supported on the constraint set. This constrained distribution arises in many applications where generated samples must strictly satisfy known physical constraints. We pose the problem in the framework of the Langevin Monte Carlo algorithm [18], a fundamental method that underlies many modern generative modeling frameworks. The application to deep generative models is discussed in Section 4.4.

Langevin Monte Carlo Sampling algorithms aim to generate random variables following a target distribution assumed to admit a density of the form $p(x) := e^{-f(x)}/Z$ on \mathbb{R}^d , where f(x) is called the potential function and is assumed to be differentiable. The Markov chain Monte Carlo methods design iterative algorithms producing samples (x_t) whose distribution q_t converges to p. Among them, the Langevin Monte Carlo algorithm plays a central role. It requires access to the gradient of the potential $\nabla f(x)$, also called the score function [19], and performs noisy gradient descent updates

$$x_{t+1} = x_t - \tau \nabla f(x_t) + \sqrt{2\tau} w_t, \qquad w_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d).$$
(2.1)

where τ is the step size. Under standard assumptions, the stationary distribution of the Markov chain converges to p [13].

Constrained target distribution We now consider the case where the samples are known to satisfy hard constraints at sampling time, in the form of a bounded measurable set $C \subset \mathbb{R}^d$. Our objective is then to sample from the conditional density supported on C:

$$p_{\mathcal{C}}(x) := \frac{1}{Z_{\mathcal{C}}} e^{-f(x)} \mathbb{1}_{\mathcal{C}}(x), \quad \forall x \in \mathbb{R}^d,$$
(2.2)

with $\mathbb{1}_{\mathcal{C}}$ the indicator function of \mathcal{C} and $Z_{\mathcal{C}}$ is a normalizing constant. This conditional distribution can be rewritten using a modified potential: $p_{\mathcal{C}}(x) := e^{-f_{\mathcal{C}}(x)}/Z_{\mathcal{C}}$, with the constrained potential $f_{\mathcal{C}}(x) := f(x) + \chi_{\mathcal{C}}(x)$, defined with the characteristic function of \mathcal{C}

$$\chi_{\mathcal{C}}(x) = \begin{cases} 0 & \text{if } x \in \mathcal{C}, \\ +\infty & \text{otherwise.} \end{cases}$$
(2.3)

We do not make any assumption on the constraint set C, except that it is bounded and that p_C is well-defined. Next, we provide examples of such constraints that may occur in physical applications.

Example 2.1 [Equality constraints] In many physical settings, the constraint set C can be characterized by equality constraints of the form $C = \{x \in \mathbb{R}^d \mid h(x) = 0\}$, where $c : \mathbb{R}^d \to \mathbb{R}^m$ is a smooth function. For instance, when x describes a discretized physical field, conservation of energy can often be expressed as $h(x) = \frac{1}{2} ||x||_2^2 - E$ for a prescribed energy E, while conservation of mass corresponds to $h(x) = \sum_i x_i - M$ for a prescribed mass M.

Objective Our objective is to design a sampling algorithm that produces samples distributed according to $p_{\mathcal{C}}$ for any constraint set \mathcal{C} . It should use only access to the score function $\nabla f(x)$ of the unconstrained density, and mathematical operations related to \mathcal{C} such as constraint functions or a projection operator $P_{\mathcal{C}}$ onto \mathcal{C} . The method should operate in a "zero-shot" scenario: it should require no access to additional data and no modification of the pre-trained generative model.

Example 2.2 [Projected Langevin] A natural idea to enforce hard constraints in Langevin dynamics is to interleave each unconstrained update (2.1) with a projection onto C, leading to

$$x_{t+1} = P_{\mathcal{C}}(x_t - \tau \nabla f(x_t) + \sqrt{2\tau} w_t), \qquad w_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d).$$

$$(2.4)$$

The corresponding constrained sampling algorithm is the projected Langevin algorithm [12, 20], which is detailed in Algorithm 4 of Appendix A, with its connection to proximal methods.

When C is convex and p is log-concave, projected Langevin and projected diffusion enjoy strong theoretical guarantees [12, 14, 20]. However, for non-convex constraint sets, hard projection at each step can significantly hinder exploration: although samples satisfy the constraint, they may be biased toward subregions of C that are easier to reach [16, 21]. This motivates the need for sampling methods that gradually enforce constraints without compromising exploration.

Algorithm evaluation Evaluating the performance of constrained sampling algorithms is challenging, as p_C is generally intractable. In practice, we rely on two key performance criteria: constraint violation and bias. First, since strict constraint satisfaction is often critical in physical applications, a crucial metric is constraint violation. It measures how far generated samples deviate from the constraint set C, for example, via a distance function or a residual. Second, even when samples lie within C, they must accurately follow the conditional distribution p_C . While direct evaluation is usually infeasible, one can assess the sampling bias by comparing empirical statistics of the samples to known or estimated quantities under p_C .

3 Variational framework of sampling and duality

In the following, we review the variational structure of Langevin dynamics and Lagrangian duality introduced by [17], which will guide the development of our strictly constrained algorithm in Section 4. Importantly, the duality framework outlined in this section enforces constraints only on average, and therefore does not directly target the strictly constrained distribution p_{C} , which is the ultimate goal of our work.

Variational view of Langevin Monte Carlo Langevin Monte Carlo admits a variational interpretation as a stochastic approximation of a gradient flow in the space of probability distributions. Let q be a probability density absolutely continuous with respect to the target density p, and define the Kullback-Leibler divergence $D(q||p) := \int_{\mathbb{R}^d} q \log(q/p)$, which is a non-negative information-theoretic quantity measuring how q differs from p [22]. We define the functional

$$F(q) := D(q||p).$$
 (3.1)

Then, Langevin dynamics can be interpreted as a stochastic particle approximation of the infinitedimensional gradient flow minimizing F in the Wasserstein space of probability measures [23, 24]. Specifically, each iteration of the Langevin algorithm corresponds to an approximate gradient step with respect to F on the distribution q_t of samples x_t , driving it toward the minimizer $q_* = p$. We refer to Appendix C for more details on this formulation.

Building on this variational formalism, constraints on average can be incorporated using classical tools from convex optimization [25]. This framework is developed in [17], where both equality and inequality constraints are considered. We focus here on equality constraints for clarity. Let $\mathcal{P}_2(\mathbb{R}^d)$ denote the set of probability measures with finite second moments, and let $h : \mathbb{R}^d \to \mathbb{R}^m$ be a constraint function. The goal is to find the closest probability distribution to p in \mathcal{P}_2 , while satisfying the constraint $\mathbb{E}_q[h(x)] = 0$, that is:

$$\begin{array}{l} \underset{q \in \mathcal{P}_{2}(\mathbb{R}^{d})}{\text{minimize}} \quad F(q) \\ \text{subject to} \quad \mathbb{E}_{q}[h(x)] = 0. \end{array}$$

$$(3.2)$$

Under standard constraint qualification assumptions (e.g., Slater's condition), Problem 3.2 is convex and admits a unique minimizer q_{\star} [17]. It is however an infinite-dimensional problem. To solve it, one can use the Lagrangian and its associated dual function.

Definition 1 [Lagrangian] The Lagrangian of Problem (3.2) is defined as

$$L(q,\lambda) := F(q) + \lambda^{\top} \mathbb{E}_{q}[h(x)] \qquad \forall q \in \mathcal{P}_{2}(\mathbb{R}^{d}), \ \lambda \in \mathbb{R}^{m}.$$
(3.3)

Definition 2 [Dual function] The dual function of Problem 3.2 is defined as

$$g(\lambda) := \inf_{q \in \mathcal{P}_2(\mathbb{R}^d)} L(q, \lambda), \qquad \forall \lambda \in \mathbb{R}^m.$$
(3.4)

The dual function g is concave, and the corresponding dual problem, consisting in maximizing $g(\lambda)$, is a finite-dimensional concave maximization problem. It provides a lower bound on the primal value, as $g(\lambda) \leq F(q_*)$ for all λ . Strong duality refers to the case of equality: $\sup g(\lambda) = F(q_*)$. A key property is that the infimum in (3.4) is achieved by $p_{\lambda}(x) \propto e^{-U(x,\lambda)}$, with the Lagrangian potential $U(x,\lambda) := f(x) + \lambda^{\top} h(x)$ Thus, optimizing $g(\lambda)$ corresponds to adjusting the constraint term added to the original potential.

Proposition 1 [Attained strong duality] Suppose that strong duality holds and is attained: there exists $\lambda_{\star} \in \mathbb{R}^{m}$ such that $g(\lambda_{\star}) = F(q_{\star})$. Then q_{\star} is the unique minimizer of the unconstrained problem:

$$\underset{q \in \mathcal{P}_2(\mathbb{R}^d)}{\text{minimize}} \quad L(q, \lambda_\star).$$
(3.5)

When strong duality is attained, Proposition 1 implies that sampling from q_{\star} can be obtained by solving a finite-dimensional unconstrained problem: finding the optimal multiplier λ_{\star} and sampling from $p_{\lambda_{\star}}$. The Lagrange multiplier can be found by to the so-called dual ascent algorithm:

$$\lambda_{t+1} = \lambda_t + \eta \mathbb{E}_{q_t}[h(x)], \tag{3.6}$$

where $q_t := p_{\lambda_t}$ and η is a step size [26]. Dual ascent is detailed in Algorithm 5. If Proposition 1 applies, this algorithm converges to λ_{\star} [17].

Primal-dual sampling In practice, the expectation $\mathbb{E}_{q_t}[h(x)]$ is approximated using samples obtained via Langevin dynamics under potential $U(x, \lambda_t)$. This motivates a primal-dual algorithm: alternating between Langevin sampling and stochastic dual ascent on λ . This scheme, proposed by Chamon et al. [17], is known as primal-dual Langevin Monte Carlo and is summarized in Algorithm 1. Although the primaldual Langevin sampling algorithm has been successfully applied to constrained sampling problems, it only enforces the constraint in expectation, without any control on the variance. However, many physical applications require the samples to be strictly constrained in C, as outlined in Section 2.

Algorithm 1 Primal-dual Langevin

input potential gradient ∇f , equality constraint function h, step sizes $\tau, \eta > 0$, iteration number T, initial distribution q_0 output sample $x_T \in \mathbb{R}^d$ define $U(x, \lambda) := f(x) + \lambda^\top h(x)$ initialize $x_0 \sim q_0, \lambda_0 = 0 \in \mathbb{R}^m$ for $0 \le t \le T - 1$ do draw $w_t \sim \mathcal{N}(0, I_d)$ $x_{t+1} = x_t - \tau \nabla_x U(x_t, \lambda_t) + \sqrt{2\tau} w_t$ $\lambda_{t+1} = \lambda_t + \eta \nabla_\lambda U(x_{t+1}, \lambda_t)$ end for

4 Split augmented Langevin for strictly constrained sampling

In this section, we introduce a novel method for the constrained sampling problem. We propose a variational interpretation of the constrained distribution $p_{\mathcal{C}}$, enabling us to apply the duality framework described in Section 3. Based on this formulation, we propose a new constrained sampling algorithm, Split Augmented Langevin (SAL), which progressively approaches $p_{\mathcal{C}}$ while ensuring that the final sample strictly satisfies the constraint.

4.1 Variational formulation of constrained sampling

Our method builds upon the following variational formulation of constrained sampling. Importantly, we observe that the constrained distribution $p_{\mathcal{C}}$ can be characterized as a projection in distribution space of the unconstrained distribution p onto the set of distributions supported on \mathcal{C} .

Proposition 2 Suppose that $\mathbb{P}_p(\mathcal{C}) > 0$. Then the conditional distribution $p_{\mathcal{C}}$ is the projection of p onto the set of distributions supported on \mathcal{C} :

$$p_{\mathcal{C}} = \underset{q \in \mathcal{P}_{2}(\mathbb{R}^{d})}{\operatorname{argmin}} \quad D(q \| p)$$
subject to $\mathbb{P}_{q}(\mathcal{C}) = 1.$

$$(4.1)$$

This projection problem is a particular case of I-projection [27], a classical object in information theory. At first glance, this problem might seem amenable to the variational framework introduced in Section 3, by expressing the support constraint as an expectation constraint $\mathbb{E}_q[\mathbb{1}_C(x) - 1] = 0$. However, the following result shows that strong duality is not attained.

Proposition 3 Strong duality holds for Problem (4.1), but is only attained for an infinite Lagrange multiplier:

$$\forall \lambda \in \mathbb{R}, \ g(\lambda) < F(q_{\star}), \quad \text{and} \quad g(\lambda) \xrightarrow{\lambda \to -\infty} F(q_{\star}).$$
(4.2)

Consequently, the dual ascent algorithm introduced in Section 3 cannot converge to the solution of (4.1). This singularity in the dual function is due to the feasible set being supported on a strict subset of \mathbb{R}^d . One possible relaxation is to allow a small violation probability, replacing $\mathbb{P}_q(\mathcal{C}) = 1$ by $\mathbb{P}_q(\mathcal{C}) \ge 1 - \delta$ for small $\delta > 0$, this however allows unphysical states and lead to poor problem conditioning. Instead, we propose a different variational relaxation, with strict constraint satisfaction.

4.2 Split Augmented Langevin

In order to relax the problem without compromising constraint satisfaction, we propose to split the variable x into a pair $(x, z) \in \mathbb{R}^d \times C$, enforcing that x and z remain close while z strictly belongs to C. We thus define a joint probability density q(x, z), with marginals q_x and q_z .

Proposition 4 [Variable splitting] Problem (4.1) is equivalent to the following problem:

$$\begin{array}{ll} \underset{q \in \mathcal{P}_{2}(\mathbb{R}^{d} \times \mathcal{C})}{\text{minimize}} & D(q_{x} \| p) \\ \text{subject to} & \mathbb{P}_{q} \left(x = z \right) = 1. \end{array}$$

$$(4.3)$$

Inspired by variable splitting techniques in optimization [28], this reformulation separates the roles of x and $z \in C$, which are respectively maximizing likelihood and enforcing the constraint. Rather than requiring x = z almost surely, we relax the condition to be satisfied in expectation, and penalize the variance. Specifically, we consider the following problem:

$$\begin{array}{ll} \underset{q \in \mathcal{P}_{2}(\mathbb{R}^{d} \times \mathcal{C})}{\text{minimize}} & D(q \| p \otimes u_{\mathcal{C}}) + \frac{\rho}{2} \mathbb{E}_{q} \left[\| x - z \|^{2} \right] \\ \text{subject to} & \mathbb{E}_{q}[x - z] = 0, \end{array}$$
(4.4)

where $u_{\mathcal{C}}$ denotes the uniform distribution on \mathcal{C} , and the penalty parameter $\rho > 0$, controls the strength the coupling. This relaxed formulation avoids the duality failure in Proposition 3 by softening the coupling constraint between x and z. Problem 4.4 then falls under the general framework of average-constrained sampling described in Section 3, with associated augmented Lagrangian potential:

$$U_{\rho}(x, z, \lambda) := f(x) + \chi_{\mathcal{C}}(z) + \lambda^{\top}(x - z) + \frac{\rho}{2} ||x - z||^2.$$
(4.5)

Stochastic primal-dual updates. Given independent Gaussian noise vectors $w_t, w'_t \sim \mathcal{N}(0, I_d)$, the stochastic updates associated with the augmented potential (4.5) are as follows:

$$x_{t+1} = x_t - \tau \left(\nabla f(x_t) + \rho(x_t - z_t + \mu_t)\right) + \sqrt{2\tau}w_t$$
(4.6a)

$$z_{t+1} = P_{\mathcal{C}}(x_{t+1} + \mu_t + \sqrt{2\tau}w_t')$$
(4.6b)

$$\mu_{t+1} = \mu_t + \eta(x_{t+1} - z_{t+1}), \tag{4.6c}$$

with rescaled multiplier $\mu := (1/\rho) \times \lambda$. We refer to this algorithm as Split Augmented Langevin, or SAL, and detail it in Algorithm 2. Appendix A gives a detailed derivation. The algorithm returns $z_T \in C$, and therefore strictly enforces the constraint.

Connection with optimization algorithms The update formulas (4.6) are similar to the Alternating Direction Method of Multipliers (ADMM) [29], widely used in constrained optimization. Here, the variables x and z play the role of the primal variables in ADMM and λ the dual, and the stochastic augmented potential (4.5) plays the role of an augmented Lagrangian. Our sampling can thus be seen a stochastic noisy analog of ADMM in sample space \mathbb{R}^d , just like Langevin Monte Carlo is an analog to the gradient descent minimization algorithm. However, it is impor-

input potential gradient ∇f , projection $P_{\mathcal{C}}$, step sizes $\tau, \eta > 0$, regularization $\rho > 0$, iteration number T, initial distribution q_0 output sample $z_T \in \mathcal{C}$ initialize $x_0 \sim q_0, z_0 = P_{\mathcal{C}}(x_0), \ \mu_0 = 0 \in \mathbb{R}^d$ for $0 \le t \le T - 1$ do draw $w_t, w'_t \sim \mathcal{N}(0, I_d)$ $x_{t+1} = x_t - \tau \nabla f(x_t) - \tau \rho(x_t - z_t + \mu_t) + \sqrt{2\tau} w_t$ $z_{t+1} = P_{\mathcal{C}}(x_{t+1} + \mu_t + \sqrt{2\tau} w'_t)$ $\mu_{t+1} = \mu_t + \eta(x_{t+1} - z_{t+1})$ end for

tant to note that our algorithm differs from the ADMM algorithm applied directly to Problem (4.1) in distribution space $\mathcal{P}_2(\mathbb{R}^d)$: our method operates on coupled samples rather than on probability distributions.

4.3 Convergence analysis

We now provide theoretical support for the proposed scheme. Proofs can be found in Appendix B. First, we prove that strong duality holds and is attained for the relaxed problem, thus ensuring the convergence of the dual ascent algorithm.

Proposition 5 [Attained duality] Strong duality holds and is attained for Problem (4.4).

Corollary 1 [Convergence guarantee] The dual ascent algorithm converges for Problem (4.4).

Corollary 1 guarantees that our relaxation leads to a well-behaved iterative algorithm. Moreover, the relaxed problem recovers the original projection in the limit of infinite coupling.

Proposition 6 [Recovery of the projection] Let q^{ρ} denote the solution to (4.4). Then

$$q_x^{\rho}, q_z^{\rho} \xrightarrow[\rho \to +\infty]{law} p_{\mathcal{C}}.$$
(4.7)

Thus, larger values of ρ bring the samples closer to the constrained distribution, while smaller values encourage exploration.

These results support SAL as an efficient and principled method for sampling from constrained distributions.

4.4 Practical implementation and deep generative models

Implementation in diffusion models Our proposed algorithm is a constrained variant of Langevin Monte Carlo, which plays a central role in many generative frameworks, including energy-based models, and diffusion models [8–11]. The split-augmented update (4.6) can be used as a drop-in replacement for standard Langevin steps, without modifying other components of the sampler such as the noise schedule or predictor architecture. This makes constraint enforcement simple and modular. Crucially, we leverage the connection between Langevin dynamics and diffusion models [1] to propose SAL as a zero-shot constrained sampling algorithm for pre-trained diffusion models. This parallel has already been exploited by Christopher et al. [14] to introduce projected diffusion models. Appendix D discusses this connection in more detail.

Constraint satisfaction Our algorithm applies to arbitrary constraint sets, provided that an approximate projection operator is available. In the case of multiple constraints, the projection step can be approximated using alternating projections [30]. Importantly, it does not require a differentiable expression of the constraints. The coupling parameter ρ , controlling the trade-off between convergence and exploration, can be tuned or progressively increased along the diffusion process to tighten constraint enforcement over time.

Computational cost In physical applications, the computational burden of learning-based methods is critical, as their main goal is to accelerate expensive simulations. Compared to unconstrained diffusion, our method adds the cost of a projection operation at each step, as does projected diffusion. The cost of this step depends on the constraint set, and generally amounts to solving a (possibly non-convex) constrained optimization problem. Nonetheless, efficient numerical methods such as augmented Lagrangian algorithms can be used to solve the projection step, and are amenable to parallelization [29, 31].

5 Application to physics-preserving generative modeling

We evaluate SAL on three scientific generative modeling tasks where physical constraints play a critical role. Our code is available at https://github.com/MB-29/constrained-sampling. We apply SAL to diffusion models as described in 4.4.

Baselines Our sampling algorithm is compared with the unconstrained Langevin algorithm and the projected Langevin algorithm and their diffusion analogs. Additionally, when the constraints are defined by equalities, we implement primal-dual Langevin of Chamon et al. [17] (Algorithm 1). All methods share the same score function, and differ only in how they incorporate constraints.

5.1 Bimodal field generation with prescribed energy

We first validate our method in a controlled setting where the constrained distribution $p_{\mathcal{C}}$ can be accurately estimated. We consider a two-dimensional field, representing for instance a fluid (see Figure 1), discretized on a 100×100 grid. A generative model samples from the equilibrium distribution p. A key macroscopic quantity is the kinetic energy, $\frac{1}{2}||x||_2^2$, which often remains conserved and is known in advance in physical prediction tasks. The task is to sample from the conditional distribution under a fixed energy $\frac{1}{2}||x||_2^2 = E$, a non-convex constraint (see Example 2.1).

Experimental setup The distribution p is bimodal in Fourier space, with asymmetric modes on the first Fourier coefficient: one positive and concentrated, the other negative and wider, allowing higher energy. The unconstrained distribution is analytically known, and is sampled with the Langevin Monte Carlo algorithm, and p_C is estimated via rejection sampling. Although simple enough for exact comparison, the bimodal nature of p makes the exploration challenging. We condition on a high energy level, only achievable via the negative mode. As the positive mode cannot satisfy the energy constraint, the correct conditional distribution concentrates on the negative mode, and we can easily compare it to the generated samples. For each method, 1000 independent chains are run and the last iterate is collected. We compute histograms of the first Fourier coefficient for evaluation.



Figure 1: Snapshot of the sampled field.



Figure 2: Empirical histograms for the first mode.

Results Figure 2 shows the results. Only SAL matches p_C closely. Projected Langevin satisfies the constraint exactly but fails to explore, yielding many samples in the wrong mode. This example shows that closest distribution satisfies the constraint on average, sampled by the primal-dual algorithm, does not match p_C . These results demonstrate that SAL enforces hard constraints while retaining enough exploration to correctly sample the conditional distribution.



Figure 3: Left Data assimilation sampled states and reanalysis. The black crosses represent the observations. **Right** Averaged relative error.

5.2 Physically-constrained data assimilation

Data assimilation, a central problem in geophysics, aims to estimate the state of a dynamical system from sparse, noisy observations using prior knowledge. Recent work applies deep generative architectures to this task [32–35], but these models do not enforce physical invariants, such as energy or mass conservation, which are essential for physical plausibility in long-term forecasting. We study physically-constrained generative models for data assimilation on the Burgers equation, a 1D reduction of the Navier-Stokes equations with conserved mass and energy that exhibits rich dynamics and complex multiscale behaviors similar to turbulence [36]. Appendix E.1 gives additional background.

Experimental setup We perform cyclic data assimilation on the Burgers equation discretized on a 200-point spatial grid. The ground truth trajectory evolves from a random initial condition over a time horizon H = 8. Observations are sparse: the system is observed at 10 equally spaced times, with 4 noisy spatial measurements at fixed, evenly spaced locations. Each method runs for 5 cycles per trajectory, producing a predicted trajectory that can be compared to the ground truth. The first baseline is 3D-Var [37], which estimates the state with a Gaussian posterior. We compare 3D-Var to deep generative models by training a diffusion model offline on a dataset of trajtories, without any conditioning. At sampling time, diffusion is combined with the Gaussian posterior, which conditions sampling to the available information. For each cycle, the analysis is computed as the average of 5 diffusion posterior samples. The experiment is repeated on 50 independent trajectories. We compute the average mean squared error with respect to the ground truth in the state space, in ℓ_2 norm, and in the constraint space, where the constraint violation error is the sample average of $h^2(x)$ for constraint function h. All methods share the same forecast model.

Results Figure 3 shows selected assimilated states and averaged error curves. In this under-observed setting, the diffusion prior helps to reconstruct the structure of complex states better than the Gaussian prior, especially for longer times, where the system shows a stiffer structure. However, unconstrained diffusion drifts away from the true trajectory, with significant deviations in both mass and energy. The projected diffusion method strictly enforces constraints but introduces high-frequency artifacts, leading to physically implausible states. Our algorithm SAL achieves the best compromise: it respects conservation and guides sampling toward physically plausible states, resulting in significantly lower estimation error. These results highlight the potential of constrained generative modeling for robust data assimilation in physical systems.

5.3 Constrained trajectory priors for feasibility problems in optimal control

As a final application, we evaluate SAL on a feasibility problem in optimal control: find trajectories that satisfy both system dynamics and non-convex obstacle avoidance constraints. These problems are hard due to the non-convexity of obstacle regions. We consider a linear dynamical system with state y(s) and control u(s), with s the physical time, and define a trajectory as $x := (y(s), u(s))_s$. Dynamics are encoded via the linear constraint set

$$\mathcal{C}_{d} := \{ x \mid y(0) = 0, \ \ddot{y}(s) = ay(s) + bu(s), \ |u(s)| \le u_{\max} \},$$
(5.1)

while obstacle constraints define $C_0 := \{x \mid y(s_i) \notin O_i\}$ for disjoint times s_i and obstacle regions O_i . The goal is to find trajectories in the intersection $C_d \cap C_0$. For this task, ADMM [38] is a classical solver alternating projections onto C_d and C_o . But when C_o is non-convex, its convergence can be compromised. Instead, we propose to guide ADMM with samples from a generative prior: a diffusion model trained on dynamics-respecting trajectories, with constraints enforced at sampling. This approach has seen promising results in control and robotics [31, 39].

Experimental setup We consider a 1D linear system subject to a linear restoring force, controlled in acceleration. A diffusion model is trained on a dataset of obstacle-free trajectories, obtained with a variety of random periodic excitations. At test time, a non-convex obstacle is introduced. The corresponding constraint is imposed during sampling. In order to avoid the obstacle, the algorithm needs to find a swinging trajectory, taking momentum then switching direction. Each sampled trajectory is then used to initialize ADMM, and we record the fraction of samples for which a feasible solution is found.

Results Figure 4 shows some sampled trajectories and success rates as the obstacle sizes r increases, computed over 1000 samples. Figure 4 shows example trajectories and success rates as obstacle difficulty (parameter r) increases. Without constraint-aware sampling, ADMM often fails. Projected diffusion avoids obstacles but distorts dynamics, producing unrealistic paths. SAL balances both aspects: it produces obstacle-avoiding trajectories that remain dynamically feasible, leading to significantly higher success.



Figure 4: Left Dashed lines are sampled trajectories, sampled lines are the projections onto the feasibility set. **Right** Feasibility success rates for different obstacle sizes.

6 Related work

Early work on constrained sampling adapted methods from classical optimization to Langevin dynamics, such as Projected Langevin Monte Carlo [12, 13], proximal Monte Carlo [20, 40–42], Mirrored Langevin [16, 43, 44], and penalized Langevin 45. Some of these approaches have been extended to diffusion models [14, 15]. While there methods benefit convergence guarantees in the convex case where imposing constraints does not impair exploration, they are less practical in non-convex settings, such as those encountered in our applications to physical systems. The variational formulation of Langevin Monte Carlo is leveraged by Salim and Richtarik [41] and Chamon et al. [17], where constraints are imposed only on average.

Sampling methods using variable splitting have recently been proposed for another type of problem: Bayesian posterior sampling [46] and, more recently, in plug-and-play generative solvers for inverse problems [47–49].

Physical constraints have been enforced on black-box deterministic neural networks [50, 51]. In another line of work, Cheng et al. [52] introduce a sampling algorithm that integrates a projection onto the feasibility set within a flow-matching framework, which is connected to but different from Langevin dynamics. Meunier et al. [53] propose to impose soft constraints into diffusion models for physically plausible ocean modeling.

7 Conclusion

We introduced a new sampling algorithm for constrained generative modeling, designed to enforce hard constraints while preserving exploration. Our method, Split Augmented Langevin (SAL), builds on a variational reformulation of conditional sampling and leverages primal-dual updates in a relaxed space where strict constraint satisfaction is progressively enforced. Unlike projection-based methods, which may severely distort sampling dynamics, SAL maintains fidelity to the target distribution while guaranteeing that all samples lie in the constraint set. The algorithm is modular and can be integrated into existing Langevin-based or diffusion-based samplers without retraining, and with minimal assumptions on the constraints. Through experiments on physical systems, including data assimilation and optimal control, we demonstrated that SAL improves constraint enforcement and predictive accuracy. These results open promising directions for combining generative models with

physical reasoning in scientific applications, where respecting conservation laws and feasibility constraints is crucial.

The limitations of our work include the computational cost induced by projecting the iterates, which, depending on the constraint, can further slow down generative models, and the choice of the coupling parameter ρ , which we found to have an impact on the results, as it does for ADMM.

It would be interesting to generalize the variational framework to other generative models, such as stochastic interpolants [54], and to other constraints, such as noisy observations. Another important research direction would be the finite-time convergence analysis of SAL, following the method of [17].

References

- [1] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [2] Gabriele Corso, Hannes StÃ, Bowen Jing, Regina Barzilay, Tommi Jaakkola, et al. DiffDock: Diffusion Steps, Twists, and Turns for Molecular Docking. In *International Conference on Learning Representations (ICLR 2023)*, 2023.
- [3] Ilan Price, Alvaro Sanchez-Gonzalez, Ferran Alet, Tom R Andersson, Andrew El-Kadi, Dominic Masters, Timo Ewalds, Jacklynn Stott, Shakir Mohamed, Peter Battaglia, et al. Probabilistic weather forecasting with machine learning. *Nature*, 637(8044):84–90, 2025.
- [4] Edward S Epstein and Rex J Fleming. Depicting stochastic dynamic forecasts. *Journal of Atmospheric Sciences*, 28(4):500–511, 1971.
- [5] Juan Nathaniel and Pierre Gentine. Generative emulation of chaotic dynamics with coherent prior. *arXiv preprint arXiv:2504.14264*, 2025.
- [6] Karthik Kashinath, M Mustafa, Adrian Albert, JL Wu, C Jiang, Soheil Esmaeilzadeh, Kamyar Azizzadenesheli, R Wang, Ashesh Chattopadhyay, A Singh, et al. Physics-informed machine learning: case studies for weather and climate modelling. *Philosophical Transactions of the Royal Society A*, 379(2194):20200093, 2021.
- [7] Chris Pedersen, Laure Zanna, and Joan Bruna. Thermalizer: Stable autoregressive neural emulation of spatiotemporal chaos. *arXiv preprint arXiv:2503.18731*, 2025.
- [8] Geoffrey E Hinton. Training products of experts by minimizing contrastive divergence. *Neural computation*, 14(8):1771–1800, 2002.
- [9] Yilun Du and Igor Mordatch. Implicit generation and modeling with energy based models. *Advances in neural information processing systems*, 32, 2019.
- [10] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- [11] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *International Conference on Learning Representations*, 2020.
- [12] Sebastien Bubeck, Ronen Eldan, and Joseph Lehec. Finite-time analysis of projected langevin monte carlo. *Advances in Neural Information Processing Systems*, 28, 2015.
- [13] Alain Durmus, Szymon Majewski, and Błażej Miasojedow. Analysis of Langevin Monte Carlo via convex optimization. *Journal of Machine Learning Research*, 20(73):1–46, 2019.
- [14] Jacob K Christopher, Stephen Baek, and Nando Fioretto. Constrained synthesis with projected diffusion models. Advances in Neural Information Processing Systems, 37:89307–89333, 2024.
- [15] Nic Fishman, Leo Klarner, Valentin De Bortoli, Emile Mathieu, and Michael John Hutchinson. Diffusion Models for Constrained Domains. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL https://openreview.net/forum?id=xuWTFQ4VG0. Expert Certification.
- [16] Kwangjun Ahn and Sinho Chewi. Efficient constrained sampling via the mirror-langevin algorithm. *Advances in Neural Information Processing Systems*, 34:28405–28418, 2021.
- [17] Luiz Chamon, Mohammad Reza Karimi Jaghargh, and Anna Korba. Constrained Sampling with Primal-Dual Langevin Monte Carlo. Advances in Neural Information Processing Systems, 37:29285–29323, 2024.
- [18] P. J. Rossky, J. D. Doll, and H. L. Friedman. Brownian dynamics as smart Monte Carlo simulation. *The Journal of Chemical Physics*, 69(10):4628–4633, 11 1978. ISSN 0021-9606. doi: 10.1063/1.436415. URL https://doi.org/10.1063/1.436415.

- [19] Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.
- [20] Nicolas Brosse, Alain Durmus, Éric Moulines, and Marcelo Pereyra. Sampling from a logconcave distribution with compact support with proximal Langevin Monte Carlo. In *Conference* on learning theory, pages 319–342. PMLR, 2017.
- [21] Rina Foygel Barber and Wooseok Ha. Gradient descent with non-convex constraints: local concavity determines convergence. *Information and Inference: A Journal of the IMA*, 7(4): 755–806, 2018.
- [22] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951.
- [23] Richard Jordan, David Kinderlehrer, and Felix Otto. The variational formulation of the Fokker– Planck equation. SIAM journal on mathematical analysis, 29(1):1–17, 1998.
- [24] Cédric Villani. *Topics in optimal transportation*, volume 58. American Mathematical Soc., 2021.
- [25] Dimitri P Bertsekas. *Constrained optimization and Lagrange multiplier methods*. Academic press, 2014.
- [26] Andrzej Ruszczyński. Constrained Optimization of Differentiable Functions, pages 286-342. Princeton University Press, 2006. ISBN 9780691119151. URL http://www.jstor.org/ stable/j.ctvcm4hcj.9.
- [27] Imre Csiszár. I-divergence geometry of probability distributions and minimization problems. *The annals of probability*, pages 146–158, 1975.
- [28] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011. ISSN 1935-8237. doi: 10.1561/2200000016. URL http://dx.doi.org/10.1561/2200000016.
- [29] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends*® *in Machine learning*, 3(1):1–122, 2011.
- [30] L.M. Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. USSR Computational Mathematics and Mathematical Physics, 7(3):200–217, 1967. ISSN 0041-5553. doi: https://doi. org/10.1016/0041-5553(67)90040-7. URL https://www.sciencedirect.com/science/ article/pii/0041555367900407.
- [31] Yorai Shaoul, Itamar Mishani, Shivam Vats, Jiaoyang Li, and Maxim Likhachev. Multi-Robot Motion Planning with Diffusion Models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=AUCYptvAf3.
- [32] Langwen Huang, Lukas Gianinazzi, Yuejiang Yu, Peter D. Dueben, and Torsten Hoefler. DiffDA: a diffusion model for weather-scale data assimilation. *Proceedings of the 41st International Conference on Machine Learning*, 2024.
- [33] François Rozet and Gilles Louppe. Score-based data assimilation. Advances in Neural Information Processing Systems, 36:40521–40541, 2023.
- [34] Yongquan Qu, Juan Nathaniel, Shuolin Li, and Pierre Gentine. Deep generative data assimilation in multimodal setting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 449–459, 2024.
- [35] Matthieu Blanke, Ronan Fablet, and Marc Lelarge. Neural Incremental Data Assimilation. In ICML 2024 AI for Science Workshop, 2024.

- [36] T. van Gastelen, W. Edeling, and B. Sanderse. Energy-conserving neural network for turbulence closure modeling. *Journal of Computational Physics*, 508:113003, 2024. ISSN 0021-9991. doi: https://doi.org/10.1016/j.jcp.2024.113003. URL https://www.sciencedirect.com/ science/article/pii/S0021999124002523.
- [37] Philippe Courtier, E Andersson, W Heckley, D Vasiljevic, M Hamrud, A Hollingsworth, F Rabier, M Fisher, and J Pailleux. The ecmwf implementation of three-dimensional variational assimilation (3d-var). i: Formulation. *Quarterly Journal of the Royal Meteorological Society*, 124(550):1783–1807, 1998.
- [38] Zuzana Bílková and Michal Šorel. Projection methods for finding intersection of two convex sets and their use in signal processing problems. *Electronic Imaging*, 33:1–6, 2021.
- [39] Joao Carvalho, An T Le, Mark Baierl, Dorothea Koert, and Jan Peters. Motion planning diffusion: Learning and planning of robot motions with diffusion models. In 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 1916–1923. IEEE, 2023.
- [40] Adil Salim, Dmitry Kovalev, and Peter Richtárik. Stochastic proximal Langevin algorithm: Potential splitting and nonasymptotic rates. Advances in Neural Information Processing Systems, 32, 2019.
- [41] Adil Salim and Peter Richtarik. Primal dual interpretation of the proximal stochastic gradient Langevin algorithm. Advances in Neural Information Processing Systems, 33:3786–3796, 2020.
- [42] Alain Durmus, Eric Moulines, and Marcelo Pereyra. Efficient bayesian computation by proximal Markov chain Monte Carlo: when Langevin meets Moreau. SIAM Journal on Imaging Sciences, 11(1):473–506, 2018.
- [43] Ya-Ping Hsieh, Ali Kavis, Paul Rolland, and Volkan Cevher. Mirrored Langevin dynamics. *Advances in Neural Information Processing Systems*, 31, 2018.
- [44] Louis Sharrock, Lester Mackey, and Christopher Nemeth. Learning rate free bayesian inference in constrained domains. In *NeurIPS*, 2023.
- [45] Mert Gurbuzbalaban, Yuanhan Hu, and Lingjiong Zhu. Penalized Overdamped and Underdamped Langevin Monte Carlo Algorithms for Constrained Sampling. *Journal of Machine Learning Research*, 25(263):1–67, 2024.
- [46] Maxime Vono, Nicolas Dobigeon, and Pierre Chainais. Split-and-Augmented Gibbs Sampler—Application to Large-Scale Inference Problems. *IEEE Transactions on Signal Processing*, 67(6):1648–1661, 2019. doi: 10.1109/TSP.2019.2894825.
- [47] Charles A Bouman and Gregery T Buzzard. Generative plug and play: Posterior sampling for inverse problems. In 2023 59th Annual Allerton Conference on Communication, Control, and Computing (Allerton), pages 1–7. IEEE, 2023.
- [48] Zihui Wu, Yu Sun, Yifan Chen, Bingliang Zhang, Yisong Yue, and Katherine Bouman. Principled probabilistic imaging using diffusion models as plug-and-play priors. Advances in Neural Information Processing Systems, 37:118389–118427, 2024.
- [49] Ségolène Martin, Anne Gagneux, Paul Hagemann, and Gabriele Steidl. PnP-Flow: Plug-andplay image restoration with flow matching. *arXiv preprint arXiv:2410.02423*, 2024.
- [50] Geoffrey Négiar, Michael W. Mahoney, and Aditi Krishnapriyan. Learning differentiable solvers for systems with hard constraints. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=vdv6CmGksr0.
- [51] Derek Hansen, Danielle C Maddix, Shima Alizadeh, Gaurav Gupta, and Michael W Mahoney. Learning physical models that can respect conservation laws. In *International Conference on Machine Learning*, pages 12469–12510. PMLR, 2023.
- [52] Chaoran Cheng, Boran Han, Danielle C Maddix, Abdul Fatir Ansari, Andrew Stuart, Michael W Mahoney, and Bernie Wang. Gradient-Free Generation for Hard-Constrained Systems. In *The Thirteenth International Conference on Learning Representations*, 2024.

- [53] Etienne Meunier, David Kamm, Guillaume Gachon, Redouane Lguensat, and Julie Deshayes. Learning to generate physical ocean states: Towards hybrid climate modeling. *arXiv preprint* arXiv:2502.02499, 2025.
- [54] Michael Samuel Albergo and Eric Vanden-Eijnden. Building Normalizing Flows with Stochastic Interpolants. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=li7qeBbCR1t.
- [55] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures.* Springer Science & Business Media, 2008.
- [56] Max Welling and Yee Whye Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, ICML'11, page 681–688, Madison, WI, USA, 2011. Omnipress. ISBN 9781450306195.
- [57] Diederik P Kingma, Max Welling, et al. Auto-encoding Variational Bayes, 2013.
- [58] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. Advances in neural information processing systems, 27, 2014.
- [59] Neil J Gordon, David J Salmond, and Adrian FM Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In *IEE proceedings F (radar and signal processing)*, volume 140, pages 107–113. IET, 1993.
- [60] Geir Evensen. The ensemble Kalman filter: Theoretical formulation and practical implementation. Ocean dynamics, 53:343–367, 2003.
- [61] Yoshikazu Sasaki. Some basic formalisms in numerical variational analysis. *Monthly Weather Review*, 98(12):875–883, 1970.
- [62] Andrew C Lorenc. Analysis methods for numerical weather prediction. Quarterly Journal of the Royal Meteorological Society, 112(474):1177–1194, 1986.
- [63] Toby van Gastelen, Wouter Edeling, and Benjamin Sanderse. Energy-conserving neural network for turbulence closure modeling. *Journal of Computational Physics*, 508:113003, 2024.

Acknowledgments We thank Carla Roesch, Luiz Chamon and Anna Korba for their insightful feedback on this work. The authors acknowledge funding, computing, and storage resources from the NSF Science and Technology Center (STC) Learning the Earth with Artificial Intelligence and Physics (LEAP) (Award #2019625).

Societal impact This work aims to improve the integration of physical constraints into generative modeling algorithms, with potential applications in scientific computing, forecasting, and control. By enabling the generation of physically consistent samples, our method may contribute to more reliable simulations in fields such as climate science, fluid dynamics, or robotics. We do not foresee any direct negative societal impacts associated with this research. The approach does not involve the collection or use of personal data, nor does it introduce mechanisms for surveillance or manipulation. Its intended applications are primarily in scientific and engineering domains, and we believe it promotes trustworthy modeling practices.

A Algorithms

A.1 Detailed algorithms

Algorithm 3 Langevin Monte Carlo	Algorithm 4 Projected Langevin Monte Carlo
input potential gradient ∇f , step size τ , itera- tion number T output sample x_T initialize $x_0 \sim q_0$ for $0 \le t \le T - 1$ do $w_t \sim \mathcal{N}(0, I_d)$ $x_{t+1} = x_t - \tau \nabla f(x_t) + \sqrt{2\tau}w_t$ end for	input potential $f(x)$, projection $P_{\mathcal{C}}$, step size τ , iteration number T output sample $x_T \in \mathcal{C}$ initialize $x_0 \sim q_0$ for $0 \le t \le T - 1$ do $w_t \sim \mathcal{N}(0, I_d)$ $x_{t+1} = z_t - \tau \nabla f(z_t) + \sqrt{2\tau}w_t$ $z_{t+1} = P_{\mathcal{C}}(x_{t+1})$ end for

A	Joorithm 6 Time-dependent SAL
Algorithm 5 Dual ascent	
input constraint function h , dual step size $\eta > 0$, iteration number T output sample x_T initialize $x_0 \in \mathbb{R}^d, \lambda_0 \in \mathbb{R}^m$ for $0 \le t \le T - 1$ do $q_t = \underset{q \in \mathcal{P}_2(\mathbb{R}^d)}{\operatorname{argmin}} L(q, \lambda_t).$ $\lambda_{t+1} = \lambda_t + \eta \mathbb{E}_{q_t}[h(x)]$	input time dependent potential gradient $f(x, t)$, iteration number T , time-dependent step sizes τ_t , projection P_c , step size $\eta > 0$, regularization $\rho > 0$, initial distribution q_0 output sample $z_T \in C$ initialize $x_0 \sim q_0$, $z_0 = P_c(x_0)$, $\mu_0 = 0 \in \mathbb{R}^d$ for $0 \le t \le T - 1$ do draw $w_t, w'_t \sim \mathcal{N}(0, I_d)$ $x_{t+1} = x_t - \tau_t \nabla f(x_t, t) - \tau_t \rho(x_t - z_t + \mu_t) + \sqrt{2\tau_t} w_t$ $z_{t+1} = P_c(x_{t+1} + \mu_t + \sqrt{2\tau_t} w'_t)$
end for	$\mu_{t+1} = \mu_t + \eta(x_{t+1} - z_{t+1})$ end for

A.2 Projected Langevin

Projected Langevin consists in applying Langevin dynamics to the constrained potential f_c , However, since f_c is non-smooth, its gradient is not defined. This issue can be addressed using the proximal operator:

$$\operatorname{prox}_{\varphi}(x) := \operatorname{argmin}_{z \in \mathbb{R}^d} \frac{1}{2} \|z - x\|^2 + \varphi(z).$$
(A.1)

An important case for non-smooth functions is the proximal operator of the characteristic function $\chi_{\mathcal{C}}$, which is the projection onto \mathcal{C} :

$$P_{\mathcal{C}}(x) := \operatorname{prox}_{\gamma_{\mathcal{C}}}(x). \tag{A.2}$$

When well-defined, the proximal operator generalizes the gradient step of a smooth function φ in the sense that $\operatorname{prox}_{\tau\varphi}(x) = x - \tau \nabla \varphi(x)$. Applying the proximal step associated with $\tau f_{\mathcal{C}}$ to the noisy iterate $x_t + \sqrt{2\tau}w_t$ yields the so-called Projected Langevin iteration $x_{t+1} = P_{\mathcal{C}}(x_t - \tau \nabla f(x_t) + \sqrt{2\tau}w_t)$. The corresponding constrained sampling algorithm is the Projected Langevin Algorithm [20], which we detail in Algorithm 4.



Figure 5: Projected Langevin Algorithm.

Limited exploration Figure 5 shows the exploration issue arising with Projected Langevin Monte Carlo in the case of non-convex constraints and a bi-modal distribution. Here, projecting on the constraint set $C = \{x \mid \frac{1}{2} ||x||^2 = E\}$ leads to poor exploration, as the samples are suck on the positive side of the likelihood landscape, while the only high-likelihood zone compatible with the constraint is on the other side.

A.3 Derivation of the Split-augmented sampling formulas

Recall the augmented Lagrangian potential

$$U_{\rho}(x, z, \lambda) := f(x) + \chi_{\mathcal{C}}(z) + \lambda^{\top}(x - z) + \frac{\rho}{2} \|x - z\|^2,$$
(A.3)

and let $\mu = (1/\rho)\lambda$. Taking a stoachastic gradient step with respect to x yields

$$x_{t+1} = x_t - \tau \left(\nabla f(x_t) + \rho(x_t - z_t + \mu_t)\right) + \sqrt{2\tau} w_t$$
(A.4)

Taking a stoachastic proximal step with respect to z yields

$$z_{t+1} = P_{\mathcal{C}}(x_{t+1} + \mu_t + \sqrt{2\tau}w_t')$$
(A.5)

Taking a stoachastic gradient step with respect to λ yields

$$\mu_{t+1} = \mu_t + \eta (x_{t+1} - z_{t+1}). \tag{A.6}$$

A.4 Connection with Split-and-Augmented Gibbs samplers

The constrained sampling formulas of SAL are related to the Split-and-Augmented Gibbs samplers of Vono et al. [46], which themselves are inspired by ADMM. The main difference is that, *z* represents a smooth, prior distribution in their case, while it represents hard constraints in our case. Therefore, the framework developed in [46] is different from the constrained sampling approach developed in our present work, and Split-and-Augmented Gibbs samplers cannot be applied to enforce strict constraints in deep generative models for example.

B Proofs

B.1 Proof of Proposition 1

The proof can be found in [17].

B.2 Proof of Proposition 2

Proposition Suppose that $\mathbb{P}_p(\mathcal{C}) > 0$. Then the conditional distribution $p_{\mathcal{C}}$ is the projection of p onto the set of distributions supported on \mathcal{C} :

$$p_{\mathcal{C}} = \underset{q \in \mathcal{P}_{2}(\mathbb{R}^{d})}{\operatorname{argmin}} \quad D(q \| p)$$
subject to $\mathbb{P}_{q}(\mathcal{C}) = 1.$
(B.1)

Proof. Let $q \in \mathcal{P}_2(\mathbb{R}^d)$ such that $\mathbb{P}_q(\mathcal{C}) = 1$. Then q vanishes almost everywhere out of \mathcal{C} . Hence,

$$D(q||p) = \int_{\mathcal{C}} q(x) \log \frac{q(x)}{p(x)} dx$$

=
$$\int_{\mathcal{C}} q(x) \log \left(\frac{q(x)}{p_{\mathcal{C}}(x)} \frac{Z_{\mathcal{C}}}{Z}\right) dx$$

=
$$D(q||p_{\mathcal{C}}) + \frac{Z_{\mathcal{C}}}{Z}$$
 (B.2)

where $Z_{\mathcal{C}}$ satisfies

$$1 = \int_{\mathbb{R}^d} p_{\mathcal{C}}$$

= $\frac{Z}{Z_{\mathcal{C}}} \int_{\mathcal{C}} p(x) dx$ (B.3)
= $\frac{Z}{Z_{\mathcal{C}}} \mathbb{P}_p(\mathcal{C}).$

Therefore,

$$D(q||p) = D(q||p_{\mathcal{C}}) + \mathbb{P}_{p}(\mathcal{C}).$$
(B.4)

This quantity is minimized for $q = p_{\mathcal{C}}$, and the minimal value is $\mathbb{P}_p(\mathcal{C})$.

B.3 Proof of Proposition 3

Proposition Consider the following problem:

$$p_{\mathcal{C}} = \underset{q \in \mathcal{P}_{2}(\mathbb{R}^{2})}{\operatorname{argmin}} \quad D(q||p)$$
subject to $\mathbb{P}_{q}(\mathcal{C}) = 1,$
(B.5)

and recall that F(q) = D(q||p) and

$$g(\lambda) := \inf_{q \in \mathcal{P}_2(\mathbb{R}^d)} L(q, \lambda).$$
(B.6)

Strong duality holds, but is attained only for an infinite Lagrange multiplier:

$$\forall \lambda \in \mathbb{R}, \ g(\lambda) < F(q_{\star}), \quad \text{and} \qquad g(\lambda) \xrightarrow[\lambda \to -\infty]{} F(q_{\star}).$$
(B.7)

Proof.

$$L(q,\lambda) = D(q||p) + \lambda \left(\mathbb{P}_q(x \in \mathcal{C}) - 1\right)$$

= $D(q||p) + \lambda \mathbb{E}_q[\mathbb{1}_{\mathcal{C}}(x) - 1].$ (B.8)

Let $h(x) = \mathbb{1}_{\mathcal{C}}(x) - 1$. For all $\lambda \in \mathbb{R}$, the infimum in the dual function definition, attained by

$$p_{\lambda}(x) = \frac{1}{Z_{\lambda}} e^{-f(x) - \lambda h(x)}$$

= $\frac{Z}{Z_{\lambda}} p(x) e^{-\lambda h(x)}$ (B.9)

and is equal to

 $g(\lambda) = \log \frac{Z}{Z_{\lambda}}.$ (B.10)

Note that

$$1 = \int_{\mathbb{R}^d} p_{\lambda}$$

= $\frac{Z}{Z_{\lambda}} \int_{\mathcal{C}} p(x) dx + \frac{Z}{Z_{\lambda}} e^{\lambda} \int_{\bar{\mathcal{C}}} p(x) dx$ (B.11)
= $\frac{Z}{Z_{\lambda}} \left[\mathbb{P}_p(\mathcal{C}) + e^{\lambda} (1 - \mathbb{P}_p(\mathcal{C})) \right],$

which gives

$$g(\lambda) = \log \frac{1}{\mathbb{P}_p(\mathcal{C}) + e^{\lambda}(1 - \mathbb{P}_p(\mathcal{C}))}.$$
(B.12)

This value is always strictly lower than its limit:

$$\forall \lambda, \ g(\lambda) < \log \frac{1}{\mathbb{P}_p(\mathcal{C})} = \lim_{\lambda \to +\infty} g(\lambda),$$
(B.13)

which is precisely the optimal value of Problem (4.1), attained by $q = p_{\mathcal{C}}$. Indeed,

$$D(p_{\mathcal{C}}||p) = \int_{\mathcal{C}} \frac{Z}{Z_{\mathcal{C}}} p(x) \log \frac{Z}{Z_{\mathcal{C}}} dx$$

= $\mathbb{P}_{p}(\mathcal{C}) \frac{Z}{Z_{\mathcal{C}}} \log \frac{Z}{Z_{\mathcal{C}}},$ (B.14)

where $Z_{\mathcal{C}}$ satisfies

$$1 = \int_{\mathbb{R}^d} p_{\mathcal{C}}$$

= $\frac{Z}{Z_{\mathcal{C}}} \int_{\mathcal{C}} p(x) dx$ (B.15)
= $\frac{Z}{Z_{\mathcal{C}}} \mathbb{P}_p(\mathcal{C}).$

It follows that

$$D(p_{\mathcal{C}}||p) = \log \frac{1}{\mathbb{P}_p(\mathcal{C})}.$$
(B.16)

This value is found to be the minimizer of Problem (4.1) using Gibbs' inequality.

B.4 Proof of Proposition 4

Proposition 7 [Variable splitting] Problem (4.1) is equivalent to the following problem:

$$\begin{array}{ll} \underset{q \in \mathcal{P}_{2}(\mathbb{R}^{d} \times \mathcal{C})}{\text{minimize}} & D(q_{x} \| p) \\ \text{subject to} & \mathbb{P}_{q} \left(x = z \right) = 1. \end{array}$$
(B.17)

Proof. Given q(x, z) the solution of Problem (4.3), the marginal q_x gives the solution of Problem (4.1). Given q(x) the solution of Problem (4.1), the solution of Problem (4.3) can be obtained by defining z as a copy of x.

B.5 Proof of Proposition 5

Proposition 8 [Attained duality] Strong duality holds and is attained for Problem (4.4).

Proof. In order to apply Proposition 2.2 from Chamon et al. [17], we verify the required assumption: there exists q > 0, such that $\mathbb{E}_q[x - z] = 0$ (positivity ensures constraint qualification). Such distribution can be obtained by defining q(x, z) := q(x)q(z|x), with for example q(x) a Gaussian normal density and q(z|x) a Gaussian density centered on x. Then, the aforementioned proposition can be applied and Proposition 5 follows. This result cannot be applied to Problem (4.1) because the feasibility set for q imposes that the density has zeros measure out of C, making the non-negativity constraint of the density not qualified.

B.6 Proof of Proposition 6

Recall the relaxed problem

$$\begin{array}{ll} \underset{q \in \mathcal{P}(\mathbb{R}^d \times \mathcal{C})}{\text{minimize}} & D(q || p \otimes u_{\mathcal{C}}) + \rho \mathbb{V}_q \left[x - z \right] \\ \text{subject to} & \mathbb{E}_q [x - z] = 0. \end{array}$$

$$(B.18)$$

Proposition [Problem approximation] The ρ -approximation converges to the strictly constrained problem, as

$$q_{\rho} \xrightarrow[\rho \to +\infty]{\text{law}} p_{\mathcal{C}}.$$

Proof of Proposition 6. Recall that, because strong duality is attained, the solution of (4.4) is attained by a distribution of the form

$$q_{\rho}(x,z) = \frac{1}{Z_{\lambda}} e^{-f(x)} e^{-\chi_{\mathcal{C}}(z)} e^{-\frac{\rho}{2} ||x-z||^2} e^{-\lambda^{\top}(x-z)}$$
(B.19)

Let $z \in \mathcal{C}$ and $x \neq z$ in \mathbb{R}^d . Then, $q_\rho(x, z) \xrightarrow[\rho \to +\infty]{} 0 = p_{\mathcal{C}}^2(x, z)$.

Additionally,

$$q_{\rho}(z,z) = \frac{1}{Z_{\lambda}} \mathrm{e}^{-f(z)} \tag{B.20}$$

	_	

C Variational framework for Langevin Monte Carlo

Consider the functional

$$F(q) = D(q||p) = \int q \log(q/p).$$
 (C.1)

The Wasserstein gradient flow is defined as the following differential system

$$\frac{\partial q}{\partial t} = \nabla \cdot \left(q \nabla \frac{\partial F}{\partial q} \right), \tag{C.2}$$

For functional (C.1), the differential system is found to be

$$\frac{\partial q}{\partial t} = \nabla \cdot (q \nabla f(x)) + \Delta q(x, t), \tag{C.3}$$

which is found to be the Fokker-Planck equation for the Langevin dynamics

$$\mathrm{d}x = -\nabla f(x)\mathrm{d}t + \mathrm{d}B. \tag{C.4}$$

More details can be found in [17, 23, 24, 55].

D Connection between Generative Models and Langevin Sampling

Many modern generative frameworks—from energy-based models to state-of-the-art diffusion models—rely on Langevin dynamics for sampling. In this appendix we briefly review how key classes of generative models relate to Langevin updates.

For these generative models, we will see that sampling takes the form

$$x_{t+1} = x_t - \tau_t \nabla f(x_t, t) + \sqrt{2\tau_t} w_t, \quad w_t \sim \mathcal{N}(0, I).$$
(D.1)

We interpret these steps as the discretization of a Wasserstein flow for a time-dependent functional F(x,t), where the derivation is identical as what is presented in Appendix C. We can then identically apply our constrained sampling algorithm, as a time-dependent variation of Algorithm 2, detailed in Algorithm 6. From a variational point of view, this results in framing the constrained sampling as a time-varying constrained optimization problem.

D.1 Energy-Based Models (EBMs)

An EBM defines a density

$$p(x) = \frac{1}{Z} \exp\left(-f_{\theta}(x)\right), \qquad (D.2)$$

where f_{θ} is a learned energy function. Sampling from *p* typically relies on Langevin dynamics (2.1) or stochastic gradient Langevin dynamics (SGLD) [56]. EBMs with Langevin sampling have demonstrated strong performance across a range of tasks[9], and offer distinct advantages over methods such as Variational Autoencoders (VAEs)[57] and Generative Adversarial Networks (GANs)[58]. A particularly valuable property of EBMs is their flexibility in incorporating constraints via summing up the corresponding energies. From this perspective, our algorithm, when applied to EBMs, can be interpreted as providing stronger constraint enforcement through an augmented Lagrangian potential and corresponding proximal Langevin updates—going beyond the simple addition of constraint energies.

D.2 Score-Based Generative Models

Score-based generative models aim to learn the score function $\nabla \log p_t(x)$ of a family of progressively noised data distributions $\{p_t\}_{t \in [0,T]}$, rather than modeling the data density directly. Once the score is learned—typically via denoising score matching—samples can be generated by Langevin-type updates.

Annealed Langevin Dynamics Proposed by Song and Ermon [10], this method generates samples by applying Langevin dynamics at a sequence of decreasing noise levels $\sigma_T > \cdots > \sigma_1$. A score model $s_{\theta}(x, \sigma)$ is trained to approximate the noise-dependent score $\nabla_x \log q(x; \sigma)$ of the perturbed data distribution $p(x; \sigma)$, which is obtained by convolving p(x) with a Gaussian of various noise level σ_t . Then update step takes the form

$$x_{t+1} = x_t + \tau_t \, s_\theta(x, \sigma_t) + \sqrt{2\tau_t} \, w_t, \quad w_t \sim \mathcal{N}(0, I), \tag{D.3}$$

where $\tau_t \propto \sigma_t^2$ are time-varying step sizes. The update takes the form of (D.1) with $\nabla f(x,t) = -s_{\theta}(x,\sigma_t)$. This can be seen as an unadjusted Langevin algorithm with temperature σ_t , gradually refining the sample as noise decreases. In this case our algorithm can be directly applied at each noise level to impose constraints. It is worth noting that the projected diffusion model [14] also falls into this category – a hard projection following each Langevin update within the annealed Langevin dynamics framework. Note that this covers the case where several Langevin steps are taken at fixed noise level, as in the work of Song and Ermon [10], by choosing τ_t to be constant for a number of steps t.

D.3 Diffusion Models

Denoising Diffusion Probabilistic Models (DDPM) Denoising diffusion probabilistic models (DDPMs), introduced by Ho et al. [1], define a forward process that gradually corrupts a data point y_0 by adding Gaussian noise through a fixed Markov chain:

$$q(y_t \mid y_{t-1}) = \mathcal{N}(y_t; \sqrt{1 - \beta_t \, y_{t-1}, \beta_t I}), \tag{D.4}$$

where $\beta_t \in (0, 1)$ is a small noise schedule. This leads to a closed-form expression for $q(x_t \mid x_0)$, with the following definitions:

$$\alpha_t = 1 - \beta_t, \quad \bar{\alpha}_t = \prod_{s=1}^t \alpha_s. \tag{D.5}$$

The reverse process is parameterized by a neural network $\epsilon_{\theta}(x_t, t)$, which predicts the noise component. The sampling procedure follows:

$$x_{t+1} = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t) \right) + \sigma_t w, \quad w_t \sim \mathcal{N}(0, I), \tag{D.6}$$

where σ_t is typically set to match the forward variance β_t . As noted by Ho et al. [1], this step corresponds to an Euler-Maruyama discretization of a variant of Langevin dynamics, and the learned noise predictor ϵ_{θ} implicitly estimates the score $\nabla \log p_t(x)$ up to a scaling factor. Hence, the sampling formula (D.6) takes the form (D.1) with $\tau_t = \sigma_t^2/2$ and

$$\nabla \log p_t(x_t) \approx s_\theta(x_t, t) = -\frac{1}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t).$$
(D.7)

The DDPM can be regarded as a discrete score-based model under the variance preserving stochastic differential equation (VP-SDE) interpretation [11], and thus our SAL sampling is valid for DDPM sampling.

D.4 Score-based Diffusion Models

Score-based diffusion models [11] directly learn the score function of perturbed data distributions and generate samples by simulating the reverse-time stochastic dynamics.

Forward SDE. Define a forward Itô SDE that gradually adds noise to data $x_0 \sim p_{data}$:

$$dx = a(x,t) dt + b(t) dW_t, \qquad (D.8)$$

where for the variance-preserving (VP) choice,

$$a(x,t) = -\frac{1}{2}\beta(t)x, \quad b(t) = \sqrt{\beta(t)}.$$
 (D.9)

This yields marginal distributions $p_t(x)$ that interpolate between the data and near-Gaussian noise as t increases.

Reverse SDE. The time-reversed process follows

$$dx = [a(x,t) - b(t)^2 \nabla_x \log p_t(x)] dt + b(t) dW'_t,$$
(D.10)

where W'_t is a reverse-time Wiener process. A neural network $s_{\theta}(x, t)$ is trained by score matching to approximate $\nabla_x \log p(x, t)$.

Predictor–Corrector sampling. Once the score network is trained, our SAL sampling is applicable. SAL can also be integrated seamlessly into the predictor-corrector sampling scheme proposed by Song et al. [11]. The predictor-corrector sampler interleaves:

• Predictor: an Euler-Maruyama step of the reverse SDE,

$$x_{t+1} = x_t - \tau_t \left[a(x_t, t) - b(t)^2 s_\theta(x_t, t) \right] + b(t) \sqrt{2\tau_t} w_t \quad w_t \sim \mathcal{N}(0, I_d).$$
(D.11)

• Corrector: a few steps of Langevin MCMC to refine samples,

$$x_{t+1} = x_t + \tau_t s_\theta(x_t, t) + \sqrt{2\tau_t} w_t, \quad w_t \sim \mathcal{N}(0, I_d).$$
 (D.12)

Similar to the previous sections, these formulas take the form of (D.1), with different time-varying potential gradients $\nabla f(x, t)$.

Summary Across EBMs, diffusion models, and hybrid schemes, the core sampling formula is an overdamped Langevin update, possibly annealed through noise scales. This makes our constrained sampling algorithm SAL compatible with all these approaches as a zero-shot plug-in.

E Experimental details

E.1 Data Assimilation

Context In many geophysical and engineering applications, one relies on numerical simulation to predict the time-dependent evolution of a complex system, whose state at physical time t is denoted by $x \in \mathbb{R}^d$. But these models are inherently imperfect—either because of computational constraints or incomplete knowledge of the true dynamics. When real-world observations $y \in \mathbb{R}^m$ become available (for example in digital-twin settings), we assume a statistical model of the form

$$y = h(x) + \varepsilon, \tag{E.1}$$

where $h : \mathbb{R}^d \to \mathbb{R}^m$ is an observation operator and ε is the measurement error. The imperfect simulation yields a prior forecast $b \in \mathbb{R}^d$, the background estimate, which must be adjusted using y to produce a more accurate estimate of the true state, usually referred to as the analysis, as the initial condition for the next simulation. Equivalently, one seeks samples from the posterior

$$p(x \mid b, y) \propto p(y \mid x) p(x \mid b). \tag{E.2}$$

This estimation problem is formulated sequentially for each new observation, by propagating the obtained posterior analysis with a forecast model, and repeating the process. Classically, this is achieved by one of three approaches: sequential Monte Carlo methods (e.g. particle filters [59]), ensemble-based filters (e.g. the Ensemble Kalman Filter [60]), or variational methods that solve for the MAP estimate (e.g. 3D-Var/4D-Var [61, 62]). The 3D-Var algorithm assumes that the background error distribution and observation error distribution are Gaussian,

$$x \mid b \sim \mathcal{N}(b, B), \quad \varepsilon \sim \mathcal{N}(0, P),$$
 (E.3)

then taking negative logarithm of (E.2) yields the following optimization target:

$$J(x) = \frac{1}{2} \left\| y - h(x) \right\|_{P^{-1}}^{2} + \frac{1}{2} \left\| x - b \right\|_{B^{-1}}^{2}.$$
 (E.4)

Data For simulating the Burgers equation, we implemented the same method as van Gastelen et al. [63], but we added an extra constant linear advection term. We work in Fourier space with the first 20 Fourier modes. The field evolves according to the Burgers equation for 4 time units. We generate 1,000 trajectories, with the field recorded at 10 timesteps for each trajectory, with the initial state drawn at random in Fourier space with a power-law decay of the coefficient magnitude.

Learning architecture We implemented a DDPM diffusion model, using the formalism detailed in Appendix D. The neural network involved is a fully connected network with depth 3 and width 128, using a cosine time embedding. It is trained for 200 epochs. At sampling time, 1000 diffusion steps are used.

Sampling The initial conditions are drawn at random following the same distribution of the training data

Additional results Figure 6 shows the evolution of key metrics for a data assimilation trajectory, for the various methods compared.



Figure 6: Mass conservation, energy conservation and ℓ_2 error.

E.2 Feasibility problem

Data The trajectories are discretized in time as $(y_1, \ldots, y_S, u_1, \ldots, u_S) \in \mathbb{R}^{2S}$, with S = 200 and a time interval $\Delta s = 0.01$. The dynamics constraint

$$\mathcal{C}_{d} := \{ x \mid y(0) = 0, \ddot{y}(s) = ay(s) + bu(s), |u(s)| \le u_{\max} \}$$
(E.5)

is described by a linear equality constraint, discretized into a linear system, and an inequality constraint on the control inputs. The projection on this convex constraint set is obtained by Dykstra's double projection algorithm [38], and is used within the ADMM solver. We take $a = -10, b = 1, u_{\text{max}} = 10$.

Learning architecture We implemented a DDPM diffusion model, using the formalism detailed in Appendix D. The neural network involved is a fully connected network with depth 3 and width 128, using a cosine time embedding. It is trained for 200 epochs. At sampling time, 1000 diffusion steps are used.

Sampling The obstacle is a line, and projecting onto the feasible region is performed by taking the corresponding state either directly above or directly underneath the obstacle.