

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 STRICTLY CONSTRAINED GENERATIVE MODELING VIA SPLIT AUGMENTED LANGEVIN SAMPLING

Anonymous authors

Paper under double-blind review

## 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 non-convex 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 (Ho et al., 2020), protein and material design (Corso et al., 2023), and probabilistic weather forecasting (Price et al., 2025). 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 (Epstein & Fleming, 1971; Nathaniel & Gentine, 2025). While perceptual applications mainly aim for plausibility, scientific and engineering problems require samples that obey strict physical or structural constraints, such as conservation laws or system dynamics (Kashinath et al., 2021). 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 (Pedersen et al., 2025). 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 (Du & Mordatch, 2019; Song et al., 2020), typically rely on 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 projected Langevin dynamics (Bubeck et al., 2015; Durmus et al., 2019; Christopher et al., 2024). 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 can cause the dynamics to become trapped in limited regions of the constraint set, hindering exploration and introducing significant sampling bias. Other approaches rather use a soft constraint penalty functions such as the barrier method (Fishman et al., 2023) and diffusion guidance (Ho & Salimans, 2022; Meunier et al., 2025), requiring a

054 differentiable constraint model. These methods encourage but do not enforce constraints, which is  
 055 insufficient when strict satisfaction is crucial. To our knowledge, no existing approach achieves both  
 056 strict constraint satisfaction and unbiased exploration.  
 057

058 **Contributions** Inspired by the variational formulation of Langevin dynamics and primal-dual op-  
 059 timization, we propose a novel sampling algorithm that bridges the gap between complex generative  
 060 modeling and constrained sampling, called Split Augmented Langevin (SAL). Our method enforces  
 061 hard constraints while preserving the exploration capability of Langevin dynamics. It ensures strict  
 062 constraint satisfaction and benefits from convergence guarantees via duality analysis. We show that  
 063 our approach generalizes to deep generative modeling and diffusion models. We demonstrate the ef-  
 064 fectionality of SAL on complex physically-constrained sampling tasks, including data assimilation  
 065 problems where maintaining physical invariants is key to reliable forecasting, and on non-convex  
 066 feasibility problems in optimal control.  
 067

## 068 2 PROBLEM FORMULATION OF CONSTRAINED LANGEVIN SAMPLING

070 In this section, we provide a mathematical formulation of constrained sampling: given a genera-  
 071 tive model and a constraint set, our goal is to generate samples from the conditional distribution  
 072 supported on the constraint set. Such constrained distributions arise in many applications where  
 073 samples must strictly satisfy known physical laws. We adopt the framework of the Langevin Monte  
 074 Carlo algorithm (Rossky et al., 1978), a foundation of modern generative modeling frameworks.  
 075 The application to deep generative models is discussed in Section 4.4.  
 076

077 **Langevin Monte Carlo** Consider a target distribution with density  $p(x) = e^{-f(x)}/Z$  on  $\mathbb{R}^d$ ,  
 078 where  $f(x)$  is a differentiable potential. Markov chain Monte Carlo methods design iterative algo-  
 079 rithms producing samples  $(x_t)$  whose distribution  $q_t$  converges to  $p$ . Among them, the Langevin Monte  
 080 Carlo algorithm plays a central role. It requires access to the gradient of the potential  $\nabla f(x)$ ,  
 081 also called the score function (Hyvärinen & Dayan, 2005), and performs noisy gradient descent  
 082 updates

$$083 x_{t+1} = x_t - \tau \nabla f(x_t) + \sqrt{2\tau} w_t, \quad w_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d), \quad (2.1)$$

084 where  $\tau$  is the step size. Under standard assumptions, the chain converges to  $p$  (Durmus et al., 2019).  
 085

086 **Constrained target distribution** We now consider the case where the samples are known to sat-  
 087 isfy hard constraints at sampling time, in the form of a bounded measurable set  $\mathcal{C} \subset \mathbb{R}^d$ , which  
 088 models prior information such as physical conservation laws. The conditional density supported  
 089 on  $\mathcal{C}$  is

$$090 p_{\mathcal{C}}(x) := \frac{1}{Z_{\mathcal{C}}} e^{-f(x)} \mathbb{1}_{\mathcal{C}}(x), \quad \forall x \in \mathbb{R}^d, \quad (2.2)$$

092 with  $\mathbb{1}_{\mathcal{C}}$  the indicator function of  $\mathcal{C}$  and  $Z_{\mathcal{C}}$  is a normalizing constant. Note that the conditional distri-  
 093 bution (2.2) can be rewritten using a modified potential:  $p_{\mathcal{C}}(x) := e^{-f_{\mathcal{C}}(x)}/Z_{\mathcal{C}}$ , with the constrained  
 094 potential  $f_{\mathcal{C}}(x) := f(x) + \chi_{\mathcal{C}}(x)$ , defined with the characteristic function of  $\mathcal{C}$

$$095 \chi_{\mathcal{C}}(x) := \begin{cases} 0 & \text{if } x \in \mathcal{C}, \\ 096 +\infty & \text{otherwise.} \end{cases} \quad (2.3)$$

098 We do not make any assumption on the constraint set  $\mathcal{C}$ , except that it is bounded and that  $p_{\mathcal{C}}$  is well-  
 099 defined. Next, we provide examples of such constraints that may occur in physical applications.

100 **Example 2.1** [Physical constraints] When  $x$  describes a discretized physical field, conservation of  
 101 energy  $E$  can often be expressed as the non-convex set  $\mathcal{C} = \{x \in \mathbb{R}^d \mid \|x\|_2^2 = E\}$ , while mass  
 102 conservation corresponds to  $\mathcal{C} = \{x \in \mathbb{R}^d \mid \sum_i x_i = M\}$  for a prescribed mass  $M$ .  
 103

104 **Objective** Our objective is to design a sampling algorithm that produces samples distributed ac-  
 105 cording to  $p_{\mathcal{C}}$  for any constraint set  $\mathcal{C}$ . It should use only access to the score function  $\nabla f(x)$  of the  
 106 unconstrained density, and mathematical operations related to  $\mathcal{C}$  such as constraint functions or a  
 107 projection operator  $P_{\mathcal{C}}$  onto  $\mathcal{C}$ . The method should operate in a “zero-shot” scenario, requiring no  
 retraining or additional data.

108 **Example 2.2** [Projected Langevin] A natural idea to enforce hard constraints is to project each  
 109 unconstrained update (2.1) onto  $\mathcal{C}$ , leading to  
 110

$$111 \quad x_{t+1} = P_{\mathcal{C}}(x_t - \tau \nabla f(x_t) + \sqrt{2\tau} w_t), \quad w_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d). \quad (2.4)$$

112 This projected Langevin algorithm, and its extension to diffusion models, enjoy strong theoretical  
 113 guarantees when  $\mathcal{C}$  is convex and  $p$  is log-concave (Bubeck et al., 2015). But with non-convex  
 114 constraints, repeated projection can trap the dynamics in small feasible regions, biasing explo-  
 115 ration (Barber & Ha, 2018; Ahn & Chewi, 2021). This motivates the need for sampling methods  
 116 that enforce constraints more gradually. More details on projected Langevin and its connection to  
 117 proximal methods can be found in Appendix A.

118 **Example 2.3** [Soft penalty methods] Constraints can also be enforced softly by adding a differen-  
 119 tiable cost  $c(x) \geq 0$  to the potential, penalizing samples far from  $\mathcal{C}$ , with a tunable coefficient  $\lambda \in \mathbb{R}$ :  
 120

$$121 \quad x_{t+1} = x_t - \tau(\nabla f(x_t) + \lambda \nabla c(x_t)) + \sqrt{2\tau} w_t, \quad w_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d). \quad (2.5)$$

122 This corresponds to guidance in diffusion models (Ho & Salimans, 2022; Huang et al., 2024b).  
 123 The cost function  $c(x)$  corresponds to a negative log-likelihood centered on the constraint set. Such  
 124 methods encourage constraint satisfaction but do not guarantee it, as violations are only smoothly pe-  
 125 nalized.

126 **Evaluation** Assessing the performance of constrained sampling algorithms is challenging, as  $p_{\mathcal{C}}$  is  
 127 generally intractable. In practice, we rely on two key performance criteria: constraint violation and  
 128 bias, which are both critical in physical applications. Constraint violation measures the deviation of  
 129 samples from  $\mathcal{C}$ , via a distance function or a residual for instance. Even when samples lie within  $\mathcal{C}$ ,  
 130 they must accurately follow the conditional distribution  $p_{\mathcal{C}}$  without bias. To quantify that, bias is  
 131 typically estimated by comparing sample statistics to known or approximated quantities under  $p_{\mathcal{C}}$ .  
 132

### 134 3 VARIATIONAL FRAMEWORK OF SAMPLING AND DUALITY

136 To better understand the constrained sampling problem, we formulate it as an optimization prob-  
 137 lem in the space of probability measures. In the following, we review the variational structure of  
 138 Langevin Monte Carlo and Lagrangian duality introduced by (Chamon et al., 2024), which will  
 139 guide the development of our strictly constrained algorithm in Section 4. Importantly, the dual-  
 140 ity framework outlined in this section enforces constraints only on average, and therefore does not  
 141 directly target the strictly constrained distribution  $p_{\mathcal{C}}$ , which is the ultimate goal of our work.

142 **Variational view of Langevin Monte Carlo** Langevin Monte Carlo admits a variational inter-  
 143 pretation as a gradient flow in the space of probability distributions. Let  $q$  be a density absolutely  
 144 continuous with respect to  $p$ , and define the Kullback-Leibler divergence with respect to  $p$   
 145

$$146 \quad F(q) := D(q||p) = \int_{\mathbb{R}^d} q \log(q/p), \quad (3.1)$$

148 which is a non-negative information-theoretic quantity measuring how  $q$  differs from  $p$  (Kullback  
 149 & Leibler, 1951). Langevin updates can be viewed as a stochastic particle approximation of the  
 150 gradient flow minimizing  $F$  in the Wasserstein space (Jordan et al., 1998; Villani, 2021). Each  
 151 iteration drives the law  $q_t$  of the chain  $(x_t)$  closer to the minimizer  $q_{\infty} = p$ . More details are given  
 152 in Appendix C.

154 **Average-constrained sampling** Building on this variational formalism, constraints can be incor-  
 155 porated on average using classical tools from convex optimization (Bertsekas, 2014). This frame-  
 156 work is developed in (Chamon et al., 2024), where both equality and inequality constraints are  
 157 considered. We focus here on equality constraints for clarity. Let  $\mathcal{P}_2(\mathbb{R}^d)$  denote the set of proba-  
 158 bility measures with finite second moments, and  $h : \mathbb{R}^d \rightarrow \mathbb{R}^m$  a constraint function. The closest  
 159 distribution to  $p$  in  $\mathcal{P}_2(\mathbb{R}^d)$  satisfying  $h(x) = 0$  on average solves

$$160 \quad \begin{aligned} & \underset{q \in \mathcal{P}_2(\mathbb{R}^d)}{\text{minimize}} && F(q) \\ & \text{subject to} && \mathbb{E}_q[h(x)] = 0. \end{aligned} \quad (3.2)$$

162 This convex problem admits a unique minimizer under standard assumptions (Chamon et al., 2024),  
 163 but remains infinite-dimensional. To solve it, one can use the Lagrangian and its associated dual  
 164 function.

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

$$166 \quad L(q, \lambda) := F(q) + \lambda^\top \mathbb{E}_q[h(x)] \quad \forall q \in \mathcal{P}_2(\mathbb{R}^d), \lambda \in \mathbb{R}^m. \quad (3.3)$$

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

$$169 \quad g(\lambda) := \inf_{q \in \mathcal{P}_2(\mathbb{R}^d)} L(q, \lambda), \quad \forall \lambda \in \mathbb{R}^m. \quad (3.4)$$

172 The dual function  $g$  is concave, and the corresponding dual problem, consisting in maximizing  $g(\lambda)$ ,  
 173 is a finite-dimensional concave maximization problem (Boyd & Vandenberghe, 2004). It provides  
 174 a lower bound on the primal value, as  $g(\lambda) \leq F(q_*)$  for all  $\lambda$ . A key property is that the infimum  
 175 in (3.4) is achieved by  $p_\lambda(x) \propto \exp(-U(x, \lambda))$ , with the Lagrangian potential

$$176 \quad U(x, \lambda) := f(x) + \lambda^\top h(x). \quad (3.5)$$

177 Strong duality refers to the case of equality, when  $\sup_{\lambda \in \mathbb{R}^m} g(\lambda) = F(q_*)$ .

179 **Proposition 1** [Attained strong duality] Suppose that strong duality holds and is attained: there  
 180 exists  $\lambda_* \in \mathbb{R}^m$  such that  $g(\lambda_*) = F(q_*)$ . Then,  $q_* = p_{\lambda_*}$ .

181 When strong duality is attained, Proposition 1 implies that sampling from  $q_*$  amounts to finding the  
 182 Lagrange multiplier  $\lambda_*$  by solving the finite-dimensional dual problem, and sampling from  $p_{\lambda_*}$ . The  
 183 Lagrange multiplier can be found by to the so-called dual ascent algorithm:

$$184 \quad \lambda_{t+1} = \lambda_t + \eta \mathbb{E}_{q_t}[h(x)], \quad q_t := p_{\lambda_t}, \quad (3.6)$$

185 where  $\eta$  is a step size (Ruszczyński, 2006). Dual ascent is detailed in Algorithm 5. If Proposition 1  
 186 applies, then this algorithm converges to  $\lambda_*$ .

188 **Primal-dual sampling** In practice, the expectation  $\mathbb{E}_{q_t}[h(x)]$  is approximated using samples  
 189 obtained via Langevin dynamics under potential  $U(x, \lambda_t)$ . This motivates a primal-dual algo-  
 190 rithm: alternating between Langevin sampling and  
 191 stochastic dual ascent on  $\lambda$ . This scheme, proposed  
 192 by Chamon et al. (2024), is known as primal-dual  
 193 Langevin Monte Carlo and is summarized in Algo-  
 194 rithm 1. Although the primal-dual Langevin sam-  
 195 pling has been successfully applied to constrained  
 196 sampling problems, it requires differentiable con-  
 197 straint functions, and it only enforces the constraint  
 198 in expectation, without any control on the variance.  
 199 Therefore, it does not directly target  $p_C$ . In the next  
 200 section, we address the problem of sampling while satisfying arbitrary constraints almost surely.

---

201 **Algorithm 1** Primal-dual Langevin

---

202 **input** potential gradient  $\nabla f$ , equality con-  
 203 straint function  $h$ , step sizes  $\tau, \eta > 0$ , iter-  
 204 ation number  $T$ , initial distribution  $q_0$   
 205 **output** sample  $x_T \in \mathbb{R}^d$   
 206 **define**  $U(x, \lambda) := f(x) + \lambda^\top h(x)$   
 207 **initialize**  $x_0 \sim q_0, \lambda_0 \in \mathbb{R}^m$   
 208 **for**  $0 \leq t \leq T - 1$  **do**  
 209    draw  $w_t \sim \mathcal{N}(0, I_d)$   
 210     $x_{t+1} = x_t - \tau \nabla_x U(x_t, \lambda_t) + \sqrt{2\tau} w_t$   
 211     $\lambda_{t+1} = \lambda_t + \eta \nabla_\lambda U(x_{t+1}, \lambda_t)$   
 212 **end for**

---

213 **4 SPLIT AUGMENTED LANGEVIN FOR STRICTLY CONSTRAINED SAMPLING**

215 In this section, we introduce a novel method for the constrained sampling problem. We first derive  
 216 a variational formulation of the constrained distribution  $p_C$  and, drawing on the duality framework  
 217 of Section 3, show why standard penalty-based methods fail to enforce strict constraints. Building on  
 218 this analysis, we propose Split Augmented Langevin (SAL), a constrained sampling algorithm that  
 219 provably approaches  $p_C$  while ensuring that all samples belong to  $\mathcal{C}$  through the use of a non-smooth  
 220 potential.

221 **4.1 VARIATIONAL FORMULATION OF CONSTRAINED SAMPLING**

223 Our method builds upon the following variational formulation of constrained sampling. Importantly,  
 224 we observe that the constrained distribution  $p_C$  can be characterized as the projection of the uncon-  
 225 strained distribution  $p$  onto the set of distributions supported on  $\mathcal{C}$ .

216 **Proposition 2** Suppose that  $0 < \mathbb{P}_p(\mathcal{C}) < 1$ . Then the conditional distribution  $p_{\mathcal{C}}$  is the projection  
 217 of  $p$  onto the set of distributions supported on  $\mathcal{C}$ :

$$\begin{aligned} 218 \quad p_{\mathcal{C}} &= \operatorname{argmin}_{q \in \mathcal{P}_2(\mathbb{R}^d)} D(q\|p) \\ 219 \\ 220 \quad \text{subject to } \mathbb{P}_q(\mathcal{C}) &= 1. \end{aligned} \tag{4.1}$$

222 This is a special case of I-projection (Csiszár, 1975). To solve it, one might try to apply the du-  
 223 ality framework of Section 3 to this problem by casting the support constraint as an expectation  
 224 constraint  $\mathbb{E}_q[c(x)] = 0$ , where  $c(x) \geq 0$  vanishes only on  $\mathcal{C}$ . The resulting Lagrangian potential  
 225 then exactly matches the penalty and guidance schemes of Example 2.3, thus providing a variational  
 226 interpretation of these approaches. However, we show in the following result that strong duality is  
 227 not attained, implying that such methods cannot ensure strict constraint satisfaction.

228 **Proposition 3** For Problem (4.1), strong duality is only attained for an infinite Lagrange multiplier:

$$\forall \lambda \in \mathbb{R}, \quad g(\lambda) < F(q_*) \quad \text{and} \quad g(\lambda) \xrightarrow{\lambda \rightarrow +\infty} F(q_*). \tag{4.2}$$

231 **Corollary 1** [Penalty methods] Penalty methods (2.5) cannot enforce  $\mathbb{P}_q(\mathcal{C}) = 1$ .

232 This singularity stems from the support set being a strict subset of  $\mathbb{R}^d$ . A possible relaxation is to  
 233 allow a small violation probability  $\mathbb{P}_q(\mathcal{C}) \geq 1 - \delta$  for small  $\delta > 0$ , but this allows unphysical states  
 234 and leads to poor conditioning. To overcome this limitation, we introduce a different relaxation that  
 235 preserves strict constraint satisfaction.

## 237 4.2 SPLIT AUGMENTED LANGEVIN

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

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

$$\begin{aligned} 243 \quad \operatorname{minimize}_{q \in \mathcal{P}_2(\mathbb{R}^d \times \mathcal{C})} \quad & D(q_x\|p) \\ 244 \\ 245 \quad \text{subject to } \mathbb{P}_q(x = z) &= 1. \end{aligned} \tag{4.3}$$

246 This formulation mirrors variable splitting techniques in optimization (Boyd et al., 2011a), and  
 247 separates the roles of  $x$  and  $z \in \mathcal{C}$ , which are respectively maximizing likelihood and enforcing  
 248 the constraint. Rather than requiring  $x = z$  almost surely, we relax the condition to be satisfied in  
 249 expectation, and penalize the variance. Specifically, we consider the following problem:

$$\begin{aligned} 251 \quad \operatorname{minimize}_{q \in \mathcal{P}_2(\mathbb{R}^d \times \mathcal{C})} \quad & D(q\|p \otimes u_{\mathcal{C}}) + \frac{\rho}{2} \mathbb{E}_q [\|x - z\|^2] \\ 252 \\ 253 \quad \text{subject to } \mathbb{E}_q[x - z] &= 0, \end{aligned} \tag{4.4}$$

254 where  $u_{\mathcal{C}}$  denotes the uniform distribution on  $\mathcal{C}$ , and parameter  $\rho > 0$  controls coupling strength.  
 255 This relaxed formulation avoids the duality failure in Proposition 3 by softening the coupling con-  
 256 straint between  $x$  and  $z$ . Following Section 3, we introduce the associated non-smooth augmented  
 257 Lagrangian potential

$$258 \quad U_{\rho}(x, z, \lambda) := f(x) + \chi_{\mathcal{C}}(z) + \lambda^{\top}(x - z) + \frac{\rho}{2} \|x - z\|^2. \tag{4.5}$$

259 **Stochastic proximal primal-dual updates.** To sample from this non-smooth potential, we gen-  
 260 eralize the primal-dual iterations of Chamon et al. (2024) to stochastic proximal iterations. Given  
 261 independent Gaussian noise vectors  $w_t, w'_t \sim \mathcal{N}(0, I_d)$ , the stochastic updates derived from the  
 262 augmented potential (4.5) are

$$264 \quad x_{t+1} = x_t - \tau (\nabla f(x_t) + \rho(x_t - z_t + \mu_t)) + \sqrt{2\tau} w_t \tag{4.6a}$$

$$265 \quad z_{t+1} = P_{\mathcal{C}}(z_t - \tau \rho(z_t - x_{t+1} - \mu_t) + \sqrt{2\tau} w'_t) \tag{4.6b}$$

$$266 \quad \mu_{t+1} = \mu_t + \eta(x_{t+1} - z_{t+1}), \tag{4.6c}$$

267 with rescaled multiplier  $\mu := (1/\rho) \times \lambda$ . We call this scheme Split Augmented Langevin, or SAL,  
 268 detailed in Algorithm 2. The output  $z_T \in \mathcal{C}$  strictly satisfies the constraint. AppendixA gives a  
 269 detailed derivation.

**Connection with optimization algorithms** The update formulas (4.6) resemble the Alternating Direction Method of Multipliers (ADMM) (Boyd et al., 2011b), widely used in constrained optimization. Here, the variables  $x$  and  $z$  play the role of the primal variables in ADMM and  $\lambda$  the dual, and the stochastic augmented potential (4.5) plays the role of an augmented Lagrangian. Our sampling scheme can be seen a stochastic analog of ADMM in sample space  $\mathbb{R}^d$ , just like Langevin Monte Carlo parallels gradient descent. However, it differs from ADMM applied in distribution space  $\mathcal{P}_2(\mathbb{R}^d)$ , as our method operates directly on coupled samples.

### 285 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.

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

**291 Corollary 2** [Convergence guarantee] The dual ascent algorithm converges for Problem (4.4).

293 Corollary 2 guarantees that our relaxation leads to a well-behaved iterative algorithm. Moreover, the  
294 relaxed problem recovers the original projection in the limit of infinite coupling.

**295 Proposition 6** [Recovery of the projection] Let  $q^\rho$  denote the solution to (4.4). Then

$$297 q_x^\rho, q_z^\rho \xrightarrow[\rho \rightarrow +\infty]{} p_C \quad \text{in distribution.} \quad (4.7)$$

299 Thus, larger values of  $\rho$  bring the  $x$  samples closer to  $\mathcal{C}$ , while smaller values encourage exploration.

300 These results support SAL as a principled method for sampling from constrained distributions.

### 302 4.4 PRACTICAL IMPLEMENTATION AND DEEP GENERATIVE MODELS

**304 Implementation in diffusion models** Our proposed algorithm is a constrained variant of Langevin  
305 Monte Carlo, which plays a central role in many generative frameworks (Du & Mordatch, 2019;  
306 Song & Ermon, 2019). The split-augmented update (4.6) can be used as a drop-in replacement for  
307 standard Langevin steps, without altering other sampler components, making constraint enforcement  
308 simple and modular. Leveraging the connection between Langevin dynamics and diffusion  
309 models (Ho et al., 2020), SAL provides a training-free constrained sampling algorithm for pre-trained  
310 diffusion models. This parallel has already been exploited by Christopher et al. (2024) to introduce  
311 projected diffusion models. We further extend SAL to latent diffusion and to incorporate partial  
312 signal observations, which is key in real-life applications. Details are discussed in Appendix A

**313 Constraint satisfaction** Our algorithm applies to arbitrary constraint sets, provided that a projec-  
314 tion operator (exact or approximate) is available. Unlike primal-dual Langevin and penalty methods,  
315 it does not require a differentiable constraint model, which can be challenging to derive (Laumond,  
316 1987). The coupling parameter  $\rho$  can be tuned or progressively increased along the diffusion  
317 process. This is detailed with ablation studies in Appendix E.

**319 Computational cost** For learning methods to accelerate large-scale physical simulations, effi-  
320 ciency is central. Crucially, SAL extends to latent diffusion, which is key to mitigate sampling  
321 costs. Compared to unconstrained diffusion, our method adds the cost of a projection operation at  
322 each step, as does projected diffusion. For non-convex constraints, efficient numerical methods such  
323 as augmented Lagrangian algorithms can be used to solve the projection step, and are amenable to  
parallelization (Boyd et al., 2011b; Liang et al., 2025). More details can be found in Appendix A.

---

### Algorithm 2 Split Augmented Langevin (SAL)

---

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

---

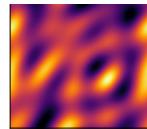
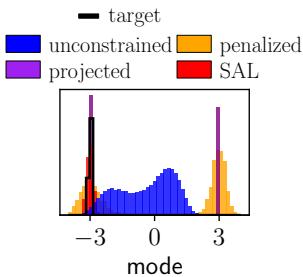
283 applied in distribution space  $\mathcal{P}_2(\mathbb{R}^d)$ , as our method operates directly on coupled samples.

324 

## 5 APPLICATION TO PHYSICS-PRESERVING GENERATIVE MODELING

325  
326 We evaluate SAL on three scientific generative modeling tasks where challenging non-convex phys-  
327 ical constraints play a critical role. We apply SAL to diffusion models as described in Section 4.4.  
328329 **Baselines** Our sampling algorithm is compared with the unconstrained Langevin algorithm, the  
330 projected Langevin algorithm, constraint penalty guidance methods, primal-dual Langevin, and their  
331 diffusion analogs (Carvalho et al., 2023; Christopher et al., 2024; Zhang et al., 2025; Zampini et al.,  
332 2025). All methods share the same score function, and differ only in how constraints are incorpo-  
333 rated. More details are given in Appendix E.334 

### 5.1 ENERGY-PRESERVING STATIONARY FIELD GENERATION

335 We first validate our method to constrained Monte Carlo sampling of a station-  
336 ary distribution. We consider a two-dimensional field, representing for instance  
337 a fluid (see Figure 1), discretized on a  $100 \times 100$  grid. The equilibrium distri-  
338 bution  $p$  is sampled using Langevin dynamics. Sampling from an equilibrium  
339 distribution is a critical problem in climate science and in molecular dynamics  
340 for example (Paquet & Viktor, 2015; Pedersen et al., 2025). A key macroscopic  
341 quantity is the kinetic energy, which often remains conserved and is known in  
342 advance in physical prediction tasks. The task is to sample from the conditional  
343 distribution under a fixed energy  $\frac{1}{2} \|x\|_2^2 = E$ , a non-convex constraint.  
344Figure 1: Sampled  
field snapshot.345 **Experimental setup** The distribution  $p$  is bimodal in Fourier  
346 space, with asymmetric modes on the first Fourier coefficient: one  
347 positive and concentrated, the other negative and wider, allowing  
348 higher energy. The unconstrained distribution is sampled with the  
349 Langevin Monte Carlo algorithm, and  $p_C$  is estimated via rejection  
350 sampling. The bimodal nature of  $p$  makes the exploration chal-  
351 lenging. We condition on a high energy level, only achievable via  
352 the negative mode. As the positive mode cannot satisfy the energy  
353 constraint, the correct conditional distribution concentrates on the  
354 negative mode, and we can easily compare it to the generated sam-  
355 ples. For each method, 1000 independent chains are run and the  
356 last iterate is collected. We compute histograms of the first Fourier  
357 coefficient for evaluation.  
358Figure 2: Empirical his-  
tograms for the first mode.359 **Results** Figure 2 shows the results. Only SAL matches  $p_C$  closely. Projected Langevin satisfies  
360 the constraint exactly but fails to explore, yielding many samples in the wrong mode. Soft constraint  
361 penalty (Zhang et al., 2025) enforce energy conservation only on average, and therefore the produced  
362 samples do not match  $p_C$ . Primal-dual Langevin yields a similar histogram. These results demon-  
363 strate that SAL enforces hard constraints while retaining enough exploration to correctly sample the  
364 conditional distribution.  
365366 

### 5.2 PHYSICALLY-CONSTRAINED DATA ASSIMILATION

367 Data assimilation, a central problem in geophysics, aims to estimate the state of a dynamical sys-  
368 tem from sparse, noisy observations using prior knowledge. Recent work applies deep generative  
369 architectures to this task (Rozet & Louppe, 2023; Qu et al., 2024), but these models do not enforce  
370 physical invariants, such as energy or mass conservation, which are essential for physical plausibility  
371 in long-term forecasting. We study physically-constrained generative models for data assimilation  
372 on the Burgers equation, a reduction of the Navier-Stokes equations with conserved mass and energy  
373 that exhibits rich dynamics and complex multiscale behaviors similar to turbulence (van Gastelen  
374 et al., 2024). Appendix E.2 gives additional background.  
375376 **Experimental setup** We perform cyclic data assimilation on the Burgers equation discretized on  
377 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

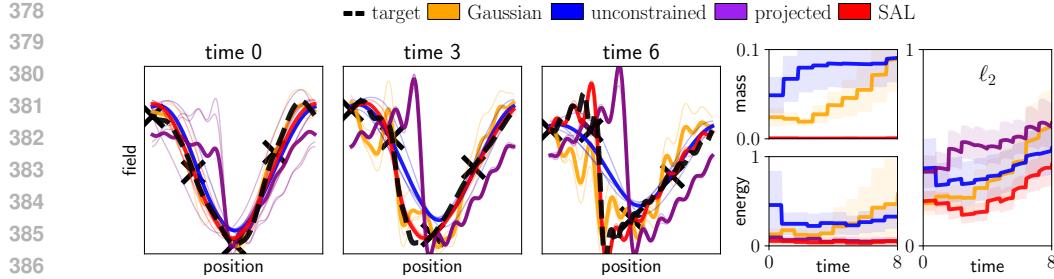


Figure 3: **Left** Data assimilation sampled states and reanalysis. The black crosses represent the observations. **Right** Averaged relative error, in terms of constraint violation and  $\ell_2$  norm.

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 (Courtier et al., 1998), which estimates the state with a Gaussian posterior. We compare 3D-Var to deep generative models by training a latent diffusion model offline on a dataset of trajectories, 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  $\ell_2$  norm, and in the constraint space, where the quadratic constraint violation error is reported. All methods share the same biased linear forecast model.

**Results** Figure 3 shows assimilated states and averaged error curves. In this under-observed setting, the diffusion prior helps to regularize 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. Projected diffusion (Christopher et al., 2024) 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 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 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 constraint set  $\mathcal{C}_d := \{x \mid \dot{y} = f(y, u), |u| \leq u_{\max}\}$ . Obstacle constraints define the potentially non-convex set  $\mathcal{C}_o := \{x \mid y(s) \notin O_i \forall s\}$ , for obstacle regions  $O_i$ . The goal is to find trajectories in the intersection  $\mathcal{C}_d \cap \mathcal{C}_o$ .

For this task, ADMM (Bílková & Šorel, 2021) is a classical solver alternating projections onto  $\mathcal{C}_d$  and  $\mathcal{C}_o$ , but its convergence can be compromised when  $\mathcal{C}_o$  is non-convex. Instead, we propose to guide ADMM with samples from a generative prior: a diffusion model trained on trajectories, with constraints enforced at sampling. This approach has seen promising results in control and robotics with diffusion penalty guidance and projected diffusion (Carvalho et al., 2023; Shaoul et al., 2025; Zampini et al., 2025), which we implement and compare with SAL.

**Experimental setup** We consider a planar quadrotor system, controlled in acceleration angle (Tedrake, 2009). A latent diffusion model is trained on a dataset of obstacle-free trajectories, obtained with a variety of random periodic excitations. At test time, non-convex obstacles are introduced. The corresponding constraint is imposed during sampling. In order to avoid the obstacles, the algorithm needs to find a swinging trajectory. Each sampled trajectory is then used to initialize ADMM, and we record the fraction of samples for which a feasible solution is found.

432 **Results** Figure 4 shows some sampled trajectories and success rates as the obstacle sizes  $r$  increases, computed over 10000 samples. Constraint penalty guidance favors obstacle avoidance, but some sampled trajectories penetrate the obstacles. Projected diffusion avoids obstacles but suffers projection bias, producing distorted and unphysical paths. Our algorithm balances both aspects: it produces obstacle-avoiding trajectories that remain dynamically feasible, leading to significantly higher success.

## 443 6 RELATED WORK

444 **Constrained Langevin Monte Carlo** Early approaches adapted optimization methods to  
 445 Langevin dynamics, including Projected Langevin Monte Carlo (Bubeck et al., 2015), proximal  
 446 Monte Carlo (Salim et al., 2019), Mirrored Langevin (Hsieh et al., 2018), and penalized  
 447 Langevin (Gurbuzbalaban et al., 2024). Extensions to diffusion models have also been ex-  
 448 plored (Fishman et al., 2023; Christopher et al., 2024; Liu et al., 2023). These methods offer con-  
 449 vergence guarantees in convex settings, where constraints do not hinder exploration, but are less  
 450 effective in non-convex physical problems. The variational formulation of Langevin sampling has  
 451 been used by Chamon et al. (2024) to enforce constraints on average.

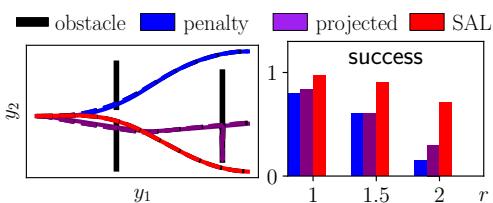
452 **Variable splitting** Variable splitting, inspired by ADMM, has been applied to Bayesian posterior  
 453 sampling (Vono et al., 2019), plug-and-play samplers for inverse problems (Bouman & Buzzard,  
 454 2023; Wu et al., 2024; Martin et al., 2024), and guided diffusion (Zhang et al., 2025). These works  
 455 apply variable splitting to a smooth maximum a posteriori optimization problem, where the aux-  
 456 illiary variable is updated by gradient descent. Crucially, our framework enforces exact constraint  
 457 satisfaction through non-smooth constraint potential, without requiring a differentiable constraint  
 458 model. Moreover, we formalize sampling as an optimization problem in density space rather than  
 459 in sample space, which is key to obtain our probabilistic sampling guarantees. Our algorithm also  
 460 extends to latent diffusion, enabling computational savings.

461 **Physically constrained neural networks** Physical constraints have also been imposed on de-  
 462 terministic neural networks (Négier et al., 2023; Hansen et al., 2023). In related sampling ap-  
 463 proaches, Cheng et al. (2024) integrate projection into flow-matching, and Meunier et al. (2025)  
 464 enforce soft constraints in diffusion models for ocean modeling. Our approach differs in targeting  
 465 strict satisfaction in a sampling framework.

## 466 7 CONCLUSION

467 We introduced Split Augmented Langevin (SAL), a new principled algorithm for constrained gen-  
 468 erative modeling that enforces hard constraints while preserving complex sampling abilities. Our  
 469 method formulates conditional sampling as a variational problem and applies primal-dual updates  
 470 in a relaxed space where strict constraints are progressively enforced. Unlike projection methods,  
 471 which can distort dynamics, or penalty methods, which may fail to enforce constraints, SAL ensures  
 472 constraint satisfaction while provably maintaining fidelity to the target distribution. The algorithm  
 473 is modular, training-free, and integrates seamlessly into Langevin samplers or diffusion models with  
 474 minimal assumptions on the constraints. Experiments on physical systems, including data assimila-  
 475 tion and optimal control, demonstrate improved constraint enforcement and predictive accuracy.  
 476 These results highlight the potential of combining generative models with physical reasoning in  
 477 scientific applications where conservation laws and feasibility constraints are essential.

478 Limitations include the computational cost of repeated projections, which may slow sampling but  
 479 can be mitigated, the choice of the coupling parameter, and the lack of a non-asymptotic anal-  
 480 ysis. Future work will extend the framework to other generative models such as stochastic in-  
 481 terpolants (Albergo & Vanden-Eijnden, 2023), and develop finite-time convergence guarantees in  
 482 Wasserstein space (Chamon et al., 2024).



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

486

**Reproducibility statement** The proofs of the new theoretical results included in this paper are  
487 available in Appendix B. The code of the proposed algorithm is available online at [anonymous.](https://anonymous.4open.science/r/constrained-sampling-F7DC/)  
488 [4open.science/r/constrained-sampling-F7DC/](https://anonymous.4open.science/r/constrained-sampling-F7DC/). Implementation details and comparison  
489 with other algorithms and diffusion models are discussed in Appendix A. Ablation studies, discus-  
490 sion about algorithm hyperparameters and experimental details are available in Appendix E.

491

492

493

494

495

496

497

498

499

500

501

502

503

504

505

506

507

508

509

510

511

512

513

514

515

516

517

518

519

520

521

522

523

524

525

526

527

528

529

530

531

532

533

534

535

536

537

538

539

540 REFERENCES  
541

542 Kwangjun Ahn and Sinho Chewi. Efficient constrained sampling via the mirror-Langevin algorithm.  
543 *Advances in Neural Information Processing Systems*, 34:28405–28418, 2021.

544 Michael Samuel Albergo and Eric Vanden-Eijnden. Building Normalizing Flows with Stochastic  
545 Interpolants. In *The Eleventh International Conference on Learning Representations*, 2023. URL  
546 <https://openreview.net/forum?id=li7qeBbCR1t>.

547 Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the*  
548 *space of probability measures*. Springer Science & Business Media, 2008.

549 Rina Foygel Barber and Wooseok Ha. Gradient descent with non-convex constraints: local concavity  
550 determines convergence. *Information and Inference: A Journal of the IMA*, 7(4):755–806, 2018.

551 Dimitri P Bertsekas. *Constrained optimization and Lagrange multiplier methods*. Academic press,  
552 2014.

553 Zuzana Bílková and Michal Šorel. Projection methods for finding intersection of two convex sets  
554 and their use in signal processing problems. *Electronic Imaging*, 33:1–6, 2021.

555 Matthieu Blanke, Ronan Fablet, and Marc Lelarge. Neural Incremental Data Assimilation. In *ICML*  
556 *2024 AI for Science Workshop*, 2024.

557 Charles A Bouman and Gregory T Buzzard. Generative plug and play: Posterior sampling for  
558 inverse problems. In *2023 59th Annual Allerton Conference on Communication, Control, and*  
559 *Computing (Allerton)*, pp. 1–7. IEEE, 2023.

560 Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed Optimization  
561 and Statistical Learning via the Alternating Direction Method of Multipliers. *Foundations and*  
562 *Trends® in Machine Learning*, 3(1):1–122, 2011a. ISSN 1935-8237. doi: 10.1561/2200000016.  
563 URL <http://dx.doi.org/10.1561/2200000016>.

564 Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. Distributed optimization  
565 and statistical learning via the alternating direction method of multipliers. *Foundations and*  
566 *Trends® in Machine learning*, 3(1):1–122, 2011b.

567 Stephen P Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

568 Nicolas Brosse, Alain Durmus, Éric Moulines, and Marcelo Pereyra. Sampling from a log-concave  
569 distribution with compact support with proximal Langevin Monte Carlo. In *Conference on learning*  
570 *theory*, pp. 319–342. PMLR, 2017.

571 Sébastien Bubeck, Ronen Eldan, and Joseph Lehec. Finite-time analysis of projected Langevin  
572 Monte Carlo. *Advances in Neural Information Processing Systems*, 28, 2015.

573 Joao Carvalho, An T Le, Mark Baierl, Dorothea Koert, and Jan Peters. Motion planning diffusion:  
574 Learning and planning of robot motions with diffusion models. In *2023 IEEE/RSJ International*  
575 *Conference on Intelligent Robots and Systems (IROS)*, pp. 1916–1923. IEEE, 2023.

576 Luiz Chamon, Mohammad Reza Karimi Jaghargh, and Anna Korba. Constrained Sampling with  
577 Primal-Dual Langevin Monte Carlo. *Advances in Neural Information Processing Systems*, 37:  
578 29285–29323, 2024.

579 Chaoran Cheng, Boran Han, Danielle C Maddix, Abdul Fatir Ansari, Andrew Stuart, Michael W  
580 Mahoney, and Bernie Wang. Gradient-Free Generation for Hard-Constrained Systems. In *The*  
581 *Thirteenth International Conference on Learning Representations*, 2024.

582 Jacob K Christopher, Stephen Baek, and Nando Fioretto. Constrained synthesis with projected  
583 diffusion models. *Advances in Neural Information Processing Systems*, 37:89307–89333, 2024.

584 Gabriele Corso, Hannes StÅ, Bowen Jing, Regina Barzilay, Tommi Jaakkola, et al. DiffDock: Dif-  
585 fusion Steps, Twists, and Turns for Molecular Docking. In *International Conference on Learning*  
586 *Representations (ICLR 2023)*, 2023.

594 Philippe Courtier, E Andersson, W Heckley, D Vasiljevic, M Hamrud, A Hollingsworth, F Rabier,  
 595 M Fisher, and J Pailleux. The ECMWF implementation of three-dimensional variational assimilation  
 596 (3D-Var). I: Formulation. *Quarterly Journal of the Royal Meteorological Society*, 124(550):  
 597 1783–1807, 1998.

598 Imre Csiszár. I-divergence geometry of probability distributions and minimization problems. *The  
 599 annals of probability*, pp. 146–158, 1975.

600 Yilun Du and Igor Mordatch. Implicit generation and modeling with energy based models. *Advances  
 601 in neural information processing systems*, 32, 2019.

602 Alain Durmus, Eric Moulines, and Marcelo Pereyra. Efficient bayesian computation by proximal  
 603 Markov chain Monte Carlo: when Langevin meets Moreau. *SIAM Journal on Imaging Sciences*,  
 604 11(1):473–506, 2018.

605 Alain Durmus, Szymon Majewski, and Błażej Miasojedow. Analysis of Langevin Monte Carlo via  
 606 convex optimization. *Journal of Machine Learning Research*, 20(73):1–46, 2019.

607 Edward S Epstein and Rex J Fleming. Depicting stochastic dynamic forecasts. *Journal of Atmo-  
 608 spheric Sciences*, 28(4):500–511, 1971.

609 Geir Evensen. The ensemble Kalman filter: Theoretical formulation and practical implementation.  
 610 *Ocean dynamics*, 53:343–367, 2003.

611 Nic Fishman, Leo Klarner, Valentin De Bortoli, Emile Mathieu, and Michael John Hutchinson.  
 612 Diffusion Models for Constrained Domains. *Transactions on Machine Learning Research*, 2023.  
 613 ISSN 2835-8856. URL <https://openreview.net/forum?id=xuWTFQ4VGO>. Expert  
 614 Certification.

615 Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,  
 616 Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information  
 617 processing systems*, 27, 2014.

618 Neil J Gordon, David J Salmond, and Adrian FM Smith. Novel approach to nonlinear/non-Gaussian  
 619 Bayesian state estimation. In *IEE proceedings F (radar and signal processing)*, volume 140, pp.  
 620 107–113. IET, 1993.

621 Mert Gurbuzbalaban, Yuanhan Hu, and Lingjiong Zhu. Penalized Overdamped and Underdamped  
 622 Langevin Monte Carlo Algorithms for Constrained Sampling. *Journal of Machine Learning Re-  
 623 search*, 25(263):1–67, 2024.

624 Derek Hansen, Danielle C Maddix, Shima Alizadeh, Gaurav Gupta, and Michael W Mahoney.  
 625 Learning physical models that can respect conservation laws. In *International Conference on  
 626 Machine Learning*, pp. 12469–12510. PMLR, 2023.

627 Geoffrey E Hinton. Training products of experts by minimizing contrastive divergence. *Neural  
 628 computation*, 14(8):1771–1800, 2002.

629 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint  
 630 arXiv:2207.12598*, 2022.

631 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in  
 632 neural information processing systems*, 33:6840–6851, 2020.

633 Ya-Ping Hsieh, Ali Kavis, Paul Rolland, and Volkan Cevher. Mirrored Langevin dynamics. *Ad-  
 634 vances in Neural Information Processing Systems*, 31, 2018.

635 Langwen Huang, Lukas Gianinazzi, Yuejiang Yu, Peter D. Dueben, and Torsten Hoefler. DiffDA: a  
 636 diffusion model for weather-scale data assimilation. *Proceedings of the 41st International Con-  
 637 ference on Machine Learning*, 2024a.

638 William Huang, Yifeng Jiang, Tom Van Wouwe, and Karen Liu. Constrained diffusion with trust  
 639 sampling. *Advances in Neural Information Processing Systems*, 37:93849–93873, 2024b.

648 Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching.  
 649 *Journal of Machine Learning Research*, 6(4), 2005.  
 650

651 Richard Jordan, David Kinderlehrer, and Felix Otto. The variational formulation of the Fokker-  
 652 Planck equation. *SIAM journal on mathematical analysis*, 29(1):1–17, 1998.  
 653

654 Karthik Kashinath, M Mustafa, Adrian Albert, JL Wu, C Jiang, Soheil Esmaeilzadeh, Kamyar Aziz-  
 655 zadenesheli, R Wang, Ashesh Chattopadhyay, A Singh, et al. Physics-informed machine learning:  
 656 case studies for weather and climate modelling. *Philosophical Transactions of the Royal Society*  
 657 A, 379(2194):20200093, 2021.  
 658

659 Diederik P Kingma, Max Welling, et al. Auto-encoding Variational Bayes, 2013.  
 660

661 Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathe-  
 662 matical statistics*, 22(1):79–86, 1951.  
 663

664 Jean-Paul Laumond. Finding Collision-Free Smooth Trajectories for a Non-Holonomic Mobile  
 665 Robot. In *IJCAI*, volume 87, pp. 1120–1123, 1987.  
 666

667 Théotime Le Hellard, Franki Nguimatsia Tiofack, Quentin Le Lidec, and Justin Carpentier. Sobolev  
 668 Diffusion Policy. 2025.  
 669

670 Jinhao Liang, Jacob K Christopher, Sven Koenig, and Ferdinando Fioretto. Multi-agent path finding  
 671 in continuous spaces with projected diffusion models. *arXiv preprint arXiv:2412.17993*, 2024.  
 672

673 Jinhao Liang, Jacob K Christopher, Sven Koenig, and Ferdinando Fioretto. Simultaneous Multi-  
 674 Robot Motion Planning with Projected Diffusion Models. *arXiv preprint arXiv:2502.03607*,  
 675 2025.  
 676

677 Guan-Horng Liu, Tianrong Chen, Evangelos Theodorou, and Molei Tao. Mirror diffusion models for  
 678 constrained and watermarked generation. *Advances in Neural Information Processing Systems*,  
 679 36:42898–42917, 2023.  
 680

681 Andrew C Lorenc. Analysis methods for numerical weather prediction. *Quarterly Journal of the  
 682 Royal Meteorological Society*, 112(474):1177–1194, 1986.  
 683

684 Ségolène Martin, Anne Gagnieux, Paul Hagemann, and Gabriele Steidl. PnP-Flow: Plug-and-play  
 685 image restoration with flow matching. *arXiv preprint arXiv:2410.02423*, 2024.  
 686

687 Etienne Meunier, David Kamm, Guillaume Gachon, Redouane Lguensat, and Julie Deshayes.  
 688 Learning to generate physical ocean states: Towards hybrid climate modeling. *arXiv preprint  
 689 arXiv:2502.02499*, 2025.  
 690

691 Juan Nathaniel and Pierre Gentine. Generative emulation of chaotic dynamics with coherent prior.  
 692 *arXiv preprint arXiv:2504.14264*, 2025.  
 693

694 Geoffrey Négiar, Michael W. Mahoney, and Aditi Krishnapriyan. Learning differentiable solvers for  
 695 systems with hard constraints. In *The Eleventh International Conference on Learning Represen-  
 696 tations*, 2023. URL <https://openreview.net/forum?id=vdv6CmGksr0>.  
 697

698 Eric Paquet and Herna L Viktor. Molecular dynamics, monte carlo simulations, and langevin dy-  
 699 namics: a computational review. *BioMed research international*, 2015(1):183918, 2015.  
 700

701 Chris Pedersen, Laure Zanna, and Joan Bruna. Thermalizer: Stable autoregressive neural emulation  
 702 of spatiotemporal chaos. *arXiv preprint arXiv:2503.18731*, 2025.  
 703

704 Ilan Price, Alvaro Sanchez-Gonzalez, Ferran Alet, Tom R Andersson, Andrew El-Kadi, Dominic  
 705 Masters, Timo Ewalds, Jacklynn Stott, Shakir Mohamed, Peter Battaglia, et al. Probabilistic  
 706 weather forecasting with machine learning. *Nature*, 637(8044):84–90, 2025.  
 707

708 Yongquan Qu, Juan Nathaniel, Shuolin Li, and Pierre Gentine. Deep generative data assimilation in  
 709 multimodal setting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern  
 710 Recognition*, pp. 449–459, 2024.

702 P. J. Rossky, J. D. Doll, and H. L. Friedman. Brownian dynamics as smart Monte Carlo simulation.  
 703 *The Journal of Chemical Physics*, 69(10):4628–4633, 11 1978. ISSN 0021-9606. doi: 10.1063/  
 704 1.436415. URL <https://doi.org/10.1063/1.436415>.

705 Litu Rout, Negin Raoof, Giannis Daras, Constantine Caramanis, Alex Dimakis, and Sanjay Shakkottai.  
 706 Solving linear inverse problems provably via posterior sampling with latent diffusion models.  
 707 *Advances in Neural Information Processing Systems*, 36:49960–49990, 2023.

708 François Rozet and Gilles Louppe. Score-based data assimilation. *Advances in Neural Information  
 709 Processing Systems*, 36:40521–40541, 2023.

710 Andrzej Ruszczyński. *Constrained Optimization of Differentiable Functions*, pp. 286–342. Prince-  
 711 ton University Press, 2006. ISBN 9780691119151. URL <http://www.jstor.org/stable/j.ctvcm4hcj.9>.

712 Adil Salim and Peter Richtarik. Primal dual interpretation of the proximal stochastic gradient  
 713 Langevin algorithm. *Advances in Neural Information Processing Systems*, 33:3786–3796, 2020.

714 Adil Salim, Dmitry Kovalev, and Peter Richtárik. Stochastic proximal Langevin algorithm: Potential  
 715 splitting and nonasymptotic rates. *Advances in Neural Information Processing Systems*, 32, 2019.

716 Yoshikazu Sasaki. Some basic formalisms in numerical variational analysis. *Monthly Weather  
 717 Review*, 98(12):875–883, 1970.

718 Yorai Shaoul, Itamar Mishani, Shivam Vats, Jiaoyang Li, and Maxim Likhachev. Multi-Robot Mo-  
 719 tion Planning with Diffusion Models. In *The Thirteenth International Conference on Learning  
 720 Representations*, 2025. URL <https://openreview.net/forum?id=AUCYptvAf3>.

721 Louis Sharrock, Lester Mackey, and Christopher Nemeth. Learning rate free Bayesian inference in  
 722 constrained domains. In *NeurIPS*, 2023.

723 Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution.  
 724 *Advances in neural information processing systems*, 32, 2019.

725 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben  
 726 Poole. Score-based generative modeling through stochastic differential equations. *International  
 727 Conference on Learning Representations*, 2020.

728 Russ Tedrake. Underactuated robotics: Learning, planning, and control for efficient and agile ma-  
 729 chines course notes for mit 6.832. *Working draft edition*, 3(4):2, 2009.

730 T. van Gastelen, W. Edeling, and B. Sanderse. Energy-conserving neural network for turbulence  
 731 closure modeling. *Journal of Computational Physics*, 508:113003, 2024. ISSN 0021-9991.  
 732 doi: <https://doi.org/10.1016/j.jcp.2024.113003>. URL <https://www.sciencedirect.com/science/article/pii/S0021999124002523>.

733 Cédric Villani. *Topics in optimal transportation*, volume 58. American Mathematical Soc., 2021.

734 Maxime Vono, Nicolas Dobigeon, and Pierre Chainais. Split-and-Augmented Gibbs Sam-  
 735 pler—Application to Large-Scale Inference Problems. *IEEE Transactions on Signal Processing*,  
 736 67(6):1648–1661, 2019. doi: 10.1109/TSP.2019.2894825.

737 Max Welling and Yee Whye Teh. Bayesian learning via stochastic gradient Langevin dynamics.  
 738 In *Proceedings of the 28th International Conference on International Conference on Machine  
 739 Learning, ICML’11*, pp. 681–688, Madison, WI, USA, 2011. Omnipress. ISBN 9781450306195.

740 Zihui Wu, Yu Sun, Yifan Chen, Bingliang Zhang, Yisong Yue, and Katherine Bouman. Princi-  
 741 paled probabilistic imaging using diffusion models as plug-and-play priors. *Advances in Neural  
 742 Information Processing Systems*, 37:118389–118427, 2024.

743 Stefano Zampini, Jacob K Christopher, Luca Oneto, Davide Anguita, and Ferdinando  
 744 Fiorotto. Training-free constrained generation with stable diffusion models. *arXiv preprint  
 745 arXiv:2502.05625*, 2025.

746 Youyuan Zhang, Zehua Liu, Zenan Li, Zhaoyu Li, James J Clark, and Xujie Si. Decoupling Training-  
 747 Free Guided Diffusion by ADMM. In *Proceedings of the Computer Vision and Pattern Recog-  
 748 nition Conference*, pp. 23292–23302, 2025.

756 **A ALGORITHMS**  
757758 **A.1 DETAILED ALGORITHMS**  
759760 **Algorithm 3** Langevin Monte Carlo  
761

---

```

input potential gradient  $\nabla f$ , step size  $\tau$ , iteration number  $T$ 
output sample  $x_T$ 
initialize  $x_0 \sim q_0$ 
for  $0 \leq t \leq T - 1$  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

```

---

771 **Algorithm 5** Dual ascent  
772

---

```

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  $0 \leq t \leq T - 1$  do
     $q_t = \operatorname{argmin}_{q \in \mathcal{P}_2(\mathbb{R}^d)} L(q, \lambda_t).$ 
     $\lambda_{t+1} = \lambda_t + \eta \mathbb{E}_{q_t}[h(x)]$ 
end for

```

---

784 **A.2 PROJECTED LANGEVIN**

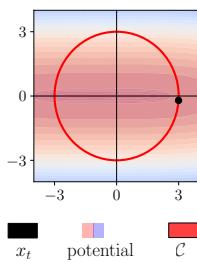
785 Projected Langevin consists in applying Langevin dynamics to the constrained potential  $f_{\mathcal{C}}$ . How-  
786 ever, since  $f_{\mathcal{C}}$  is non-smooth, its gradient is not defined. This issue can be addressed using the  
787 proximal operator:

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

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

$$P_{\mathcal{C}}(x) := \operatorname{prox}_{\chi_{\mathcal{C}}}(x). \quad (\text{A.2})$$

795 When well-defined, the proximal operator generalizes the gradient step of a smooth func-  
796 tion  $\varphi$  in the sense that  $\operatorname{prox}_{\tau\varphi}(x) = x - \tau \nabla \varphi(x)$ . Applying the proximal step associ-  
797 ated with  $\tau f_{\mathcal{C}}$  to the noisy iterate  $x_t + \sqrt{2\tau} w_t$  yields the so-called Projected Langevin iteration  
798  $x_{t+1} = P_{\mathcal{C}}(x_t - \tau \nabla f(x_t) + \sqrt{2\tau} w_t)$ . The corresponding constrained sampling algorithm is  
799 the Projected Langevin Algorithm (Brosse et al., 2017), which we detail in Algorithm 4.  
800

801   
802   
803   
804   
805   
806   
807   
808   
809 Figure 5: Projected Langevin Algorithm.

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

815     A.3 DERIVATION OF THE SPLIT-AUGMENTED SAMPLING FORMULAS

816     Recall the augmented Lagrangian potential

$$817 \quad U_\rho(x, z, \lambda) := f(x) + \chi_C(z) + \lambda^\top(x - z) + \frac{\rho}{2}\|x - z\|^2, \quad (\text{A.3})$$

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

$$819 \quad x_{t+1} = x_t - \tau(\nabla f(x_t) + \rho(x_t - z_t + \mu_t)) + \sqrt{2\tau}w_t \quad (\text{A.4})$$

820     Taking a stochastic proximal step with respect to  $z$  yields

$$821 \quad z_{t+1} = P_C(z_t - \tau\rho(z_t - x_{t+1} - \mu_t) + \sqrt{2\tau}w'_t) \quad (\text{A.5})$$

822     Taking a stochastic gradient step with respect to  $\lambda$  yields

$$823 \quad \mu_{t+1} = \mu_t + \eta(x_{t+1} - z_{t+1}). \quad (\text{A.6})$$

824     A.4 EXTENSION OF SAL TO DEEP GENERATIVE MODELS

825     Many modern generative frameworks—from energy-based models to state-of-the-art diffusion  
 826     models rely on Langevin dynamics for sampling (Hinton, 2002; Du & Mordatch, 2019; Song & Ermon,  
 827     2019; Song et al., 2020). In Appendix D, we review how key classes of generative models relate to  
 828     Langevin updates.

829     For these generative models, sampling takes the form

$$830 \quad x_{t+1} = x_t - \tau_t \nabla f(x_t, t) + \sqrt{2\tau_t}w_t, \quad w_t \sim \mathcal{N}(0, I). \quad (\text{A.7})$$

831     We interpret these steps as the discretization of a Wasserstein flow for a time-dependent functional  
 832      $F(q, t)$ , which is summarized in Appendix C. We can then identically apply our constrained  
 833     sampling algorithm, as a time-dependent variation of Algorithm 2, detailed in Algorithm 6. From  
 834     a variational point of view, this results in framing the constrained sampling as a time-varying con-  
 835     strained optimization problem.

836     A.5 EXTENSION TO LATENT DIFFUSION AND PARTIAL OBSERVATIONS

837     Our framework naturally extends to the case where the Langevin steps, or the diffusion model, are  
 838     operated in a latent space  $\mathbb{R}^k$ , which is mapped to the physical space through a decoder

$$839 \quad \varphi : \mathbb{R}^k \rightarrow \mathbb{R}^d. \quad (\text{A.8})$$

840     The augmented potential takes the form

$$841 \quad U_\rho(x, z, \lambda) := f(x) + \chi_C(z) + \lambda^\top(\varphi(x) - z) + \frac{\rho}{2}\|\varphi(x) - z\|^2, \quad (\text{A.9})$$

842     and the stochastic updates (4.6) become

$$843 \quad x_{t+1} = x_t - \tau(\nabla f(x_t) + \rho J_\varphi^\top(x_t)(\varphi(x_t) - z_t + \mu_t)) + \sqrt{2\tau}w_t \quad (\text{A.10a})$$

$$844 \quad z_{t+1} = P_C(z_t - \tau\rho(z_t - \varphi(x_{t+1}) - \mu_t) + \sqrt{2\tau}w'_t) \quad (\text{A.10b})$$

$$845 \quad \mu_{t+1} = \mu_t + \eta(\varphi(x_{t+1}) - z_{t+1}), \quad (\text{A.10c})$$

846     where  $J_\varphi$  is the decoder Jacobian.

847     When linear observations are available, in the form of a likelihood

$$848 \quad y \sim Hz + w, \quad w \sim \mathcal{N}(0, R), \quad (\text{A.11})$$

864 this additional source of information can naturally be accounted by our framework, by modifying  
 865 the augmented potential as  
 866

$$868 \quad U_\rho(x, z, \lambda) := f(x) + \chi_C(z) + \lambda^\top(\varphi(x) - z) + \frac{\rho}{2}\|\varphi(x) - z\|^2 + \frac{1}{2}\|Hz - y\|_{R^{-1}}^2, \quad (\text{A.12})$$

870 and sampling from the corresponding distribution accordingly.  
 871

## 872 A.6 CHOICE OF ALGORITHM HYPERPARAMETERS

874 The hyperparameters of our algorithm are the primal initial distribution  $q_0$ , step size  $\tau$ , the initial  
 875 dual variable  $\lambda_0$ , the dual step size  $\eta$ , and the coupling parameter  $\rho$ .

876 In practice, in diffusion models,  $q_0$  is a normal distribution, and  $\tau$  follows a prescribed schedule. We  
 877 choose  $\lambda_0 = \mu_0 = 0$  and we tune  $\eta \in [10^{-3}, 10^{-1}]$  so that we observe convergence in expectation of  
 878 the two variables  $x$  and  $z$ . The coupling parameter can be fixed or tuned during sampling. Note that,  
 879 in diffusion models, the time-varying step size directly influence the effective coupling parameter.

880 In Appendix E, we detail the hyerparameters for each experiment. We also provide ablation studies  
 881 are provided to study the influence of  $\lambda_0$  and  $\rho$  quantitatively.  
 882

## 884 A.7 COMPARISON OF CONSTRAINED SAMPLING ALGORITHM

887 **Connection with constrained Langevin Monte Carlo** Constrained sampling adapted methods  
 888 from classical optimization include Projected Langevin Monte Carlo (Bubeck et al., 2015; Durmus  
 889 et al., 2019), proximal Monte Carlo (Salim et al., 2019; Salim & Richtarik, 2020; Brosse et al., 2017;  
 890 Durmus et al., 2018), Mirrored Langevin (Hsieh et al., 2018; Ahn & Chewi, 2021; Sharrock et al.,  
 891 2023), and penalized Langevin (Gurbuzbalaban et al., 2024). These approaches are designed for  
 892 convex constraint sets, for which they enjoy theoretical guarantees. Those guarantees do not hold  
 893 for non-convex settings, such as those encountered in our applications to physical systems, where  
 894 exploration of sample space is key.

895 **Connection with Split-and-Augmented Gibbs samplers** The constrained sampling formulas of  
 896 SAL are related to the Split-and-Augmented Gibbs samplers of Vono et al. (2019), which themselves  
 897 are inspired by ADMM. The main difference is that,  $z$  represents a smooth, prior distribution in their  
 898 case, while it represents hard constraints in our case. Therefore, the framework developed in (Vono  
 899 et al., 2019) is different from the constrained sampling approach developed in our present work,  
 900 and Split-and-Augmented Gibbs samplers cannot be applied to enforce strict constraints in deep  
 901 generative models for example.

902 **Comparison with diffusion guidance** Enforcing constraints in diffusion models using penalty  
 903 and guidance methods has been proposed by Huang et al. (2024b) and Carvalho et al. (2023).  
 904 Crucially, unlike SAL, constraint penalty and guidance methods rely on a differentiable constraint  
 905 penalty function. Therefore, their sampling objective is inherently different from (4.1), as the almost-  
 906 sure constraint defines a non-smooth potential (4.5) that is not differentiable. Our approach tackles  
 907 this non-smoothness using a proximal operator and projections, thereby ensuring strict constraint  
 908 satisfaction, rather gradient steps on a smooth approached loss. In the experiments of Section 5.1,  
 909 we implement the algorithm of Zhang et al. (2025) applying guidance with variable splitting. In the  
 910 other experiments, we tried to generalize it to latent space diffusion but found the sampling algo-  
 911 rithm to fail in sampling with constraints. This is due to the difficulty of scaling the augmentation  
 912 term and the constraint penalty term to simultaneously ensure constraint satisfaction and coupling  
 913 with the diffusion model. Instead, we propagated the penalty function through the decoder.

914 **Comparison with projected diffusion** Projected diffusion can take different forms (Liang et al.,  
 915 2024; Christopher et al., 2024). Projected diffusion can be applied in a latent space Zampini et al.  
 916 (2025), which we implement in 5.3. The constraints are enforced in the latent space by solving  
 917 an optimization problem through the decoder. Our algorithm allows us to decouple these two

918 problems of sampling and of projection by relaxing the physical and the latent variable to be equal  
 919 only on average. Specifically, the iterations of Zampini et al. (2025) can be obtained as a special  
 920 case of (A.10) by setting  $\lambda_0 = 0$ ,  $\eta = 0$ , and by projecting the maximum likelihood iterate  $\varphi(x_{t+1})$ .  
 921 In our experiments, we tried to extend this algorithm to account for linear observations, but found  
 922 no simple way of propagating both a hard constraint projection and an observation likelihood to the  
 923 latent state.

#### 925 A.8 COMPUTATIONAL COST

927 To assess the computational cost of SAL, we summarize below the costs for the constraint sets used  
 928 in our experiments. We provide below a breakdown of the projection cost for different constraint  
 929 types, and compare them to the overall runtime.

930 In our experiments, we used three types of projections:

- 932 • projection onto a sphere for energy conservation,
- 933 • projection onto intervals for obstacle avoidance,
- 934 • projection onto a linear subspace for mass conservation,

936 All of these admit efficient implementations.

937 For the sphere and the interval, cost is  $\mathcal{O}(d)$  with closed-form formulas.

939 For a linear subspace,

$$940 \mathcal{C} = \{x \mid Ax + b = 0\} \quad (\text{A.13})$$

941 with  $A \in \mathbb{R}^{m \times d}$ . The projection is given by  $P_{\mathcal{C}}(x) = x + A^{\top}(AA^{\top})^{-1}(b - Ax)$ , which requires  
 942 a precomputed pseudoinverse at cost  $\mathcal{O}(d^2m)$  and a matrix-vector product at cost  $\mathcal{O}(m)$ , which  
 943 remains small compared to neural network evaluations.

944 In all tested settings, the runtime overhead from projections was small compared to the cost of score  
 945 evaluations in diffusion models. Furthermore, SAL is compatible with approximate projections,  
 946 allowing further savings. For more complex constraint sets, iterative solvers such as ADMM can  
 947 be employed with a limited number of steps, trading accuracy for speed in early iterations, where  
 948 perfect constraint enforcement is not yet required.

949 To validate this point, we conducted the following additional runtime experiment. We measure  
 950 the average wall-clock time for the different sampling algorithms in the data assimilation problem,  
 951 where each sampled state is projected on the intersection of 2 constraint sets: one for mass and one  
 952 for energy. All times are in seconds per  $10^6$  sampling steps, measured on an Apple M1 setup.

954 Experiment	955 unconstrained Langevin	956 primal-dual Langevin	957 projected Langevin	958 SAL
955 Fluid generation	956 0.37	957 0.41	958 0.43	959 0.45
956 Data assimilation	957 1.27	958 N/A	959 1.34	960 1.36

957 Table 1: Comparison of computational times.

959 We also note that diffusion guidance can be computationally costly as it can require multiple penalty  
 960 gradient steps per sampling step to enforce constraints. In the experiment 5.3, we found this method  
 961 to be of the order of 3 times slower than SAL and projected diffusion.

963 In summary, SAL has comparable runtime to projected Langevin. For a number of usecases, its ad-  
 964 dditional cost is modest, especially in the context of deep generative models where the computational  
 965 budget is dominated by score evaluations. This cost can also be adjusted in practice, by computing  
 966 approximate projections in the early steps of sampling.

972 **B PROOFS**973 **B.1 PROOF OF PROPOSITION 1**974 The proof can be found in the work of Chamon et al. (2024).  
975976 **B.2 PROOF OF PROPOSITION 2**  
977978 **Proposition** Suppose that  $\mathbb{P}_p(\mathcal{C}) > 0$ . Then the conditional distribution  $p_{\mathcal{C}}$  is the projection of  $p$   
979 onto the set of distributions supported on  $\mathcal{C}$ :  
980

981 
$$p_{\mathcal{C}} = \underset{q \in \mathcal{P}_2(\mathbb{R}^d)}{\operatorname{argmin}} D(q\|p) \quad (\text{B.1})$$
  
982 subject to  $\mathbb{P}_q(\mathcal{C}) = 1$ .  
983

984 *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,  
985

986 
$$\begin{aligned} D(q\|p) &= \int_{\mathcal{C}} q(x) \log \frac{q(x)}{p(x)} dx \\ &= \int_{\mathcal{C}} q(x) \log \left( \frac{q(x)}{p_{\mathcal{C}}(x)} \frac{Z_{\mathcal{C}}}{Z} \right) dx \\ &= D(q\|p_{\mathcal{C}}) + \frac{Z_{\mathcal{C}}}{Z} \end{aligned} \quad (\text{B.2})$$
  
987

988 where  $Z_{\mathcal{C}}$  satisfies  
989

990 
$$\begin{aligned} 1 &= \int_{\mathbb{R}^d} p_{\mathcal{C}} \\ &= \frac{Z}{Z_{\mathcal{C}}} \int_{\mathcal{C}} p(x) dx \\ &= \frac{Z}{Z_{\mathcal{C}}} \mathbb{P}_p(\mathcal{C}). \end{aligned} \quad (\text{B.3})$$
  
991

992 Therefore,  
993

994 
$$D(q\|p) = D(q\|p_{\mathcal{C}}) + \mathbb{P}_p(\mathcal{C}). \quad (\text{B.4})$$
  
995

996 This quantity is minimized for  $q = p_{\mathcal{C}}$ , and the minimal value is  $\mathbb{P}_p(\mathcal{C})$ .  $\square$   
9971000 **B.3 PROOF OF PROPOSITION 3**  
10011002 **Proposition** Consider the following problem:  
1003

1004 
$$\begin{aligned} p_{\mathcal{C}} &= \underset{q \in \mathcal{P}_2(\mathbb{R}^d)}{\operatorname{argmin}} D(q\|p) \\ &\text{subject to } \mathbb{P}_q(\mathcal{C}) = 1, \end{aligned} \quad (\text{B.5})$$
  
1005

1006 with the constrained expressed as  $\mathbb{E}[c(x)] = 0$  for a penalty function  $c(x) \geq 0$  such that  $c(x) = 0$   
1007 only on  $\mathcal{C}$ . For example,  $c(x) = 1 - \mathbb{1}_{\mathcal{C}}(x)$ . Recall that  $F(q) = D(q\|p)$  and  
1008

1009 
$$g(\lambda) := \inf_{q \in \mathcal{P}_2(\mathbb{R}^d)} L(q, \lambda). \quad (\text{B.6})$$
  
1010

1011 Strong duality holds, but is attained only for an infinite Lagrange multiplier:  
1012

1013 
$$\forall \lambda \in \mathbb{R}, g(\lambda) < F(q_{\star}), \quad \text{and} \quad g(\lambda) \xrightarrow{\lambda \rightarrow +\infty} F(q_{\star}). \quad (\text{B.7})$$
  
1014

1015 *Proof.*  
1016

1017 
$$\begin{aligned} L(q, \lambda) &= D(q\|p) + \lambda (1 - \mathbb{P}_q(x \in \mathcal{C})) \\ &= D(q\|p) + \lambda \mathbb{E}_q[c(x)]. \end{aligned} \quad (\text{B.8})$$
  
1018

1026 For all  $\lambda \in \mathbb{R}$ , the infimum in the dual function definition is attained by  
 1027

$$\begin{aligned} 1028 \quad p_\lambda(x) &= \frac{1}{Z_\lambda} e^{-f(x)-\lambda c(x)} \\ 1029 \\ 1030 \quad &= \frac{Z}{Z_\lambda} p(x) e^{-\lambda c(x)} \end{aligned} \tag{B.9}$$

1031 and the Lagrangian evaluated at  $p_\lambda$  equals  
 1032

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

1033 To compute  $Z_\lambda$ , we note that  
 1034

$$\begin{aligned} 1035 \quad 1 &= \int_{\mathbb{R}^d} p_\lambda \\ 1036 \\ 1037 \quad &= \frac{Z}{Z_\lambda} \int_{\mathcal{C}} p(x) dx + \frac{Z}{Z_\lambda} \int_{\mathcal{C}} e^{-\lambda c(x)} p(x) dx \end{aligned} \tag{B.11}$$

1038 Let  
 1039

$$\varepsilon(\lambda) := \int_{\mathcal{C}} e^{-\lambda c(x)} p(x) dx. \tag{B.12}$$

1040 Then,  
 1041

$$1 = \frac{Z}{Z_\lambda} [\mathbb{P}_p(\mathcal{C}) + \varepsilon(\lambda)] \tag{B.13}$$

1042 By assumption, for all  $\lambda \in \mathbb{R}^d$ ,  $0 < \varepsilon(\lambda) < 1$ . Furthermore, we obtain by combining (B.10)  
 1043 and (B.13), that  
 1044

$$g(\lambda) = \log \frac{1}{\mathbb{P}_p(\mathcal{C}) + \varepsilon(\lambda)}. \tag{B.14}$$

1045 This value is always strictly lower than its limit:  
 1046

$$\forall \lambda, g(\lambda) < \log \frac{1}{\mathbb{P}_p(\mathcal{C})} = \lim_{\lambda \rightarrow +\infty} g(\lambda), \tag{B.15}$$

1047 which is precisely the optimal value of Problem (4.1), attained by  $q = p_{\mathcal{C}}$ . Indeed,  
 1048

$$\begin{aligned} 1049 \quad D(p_{\mathcal{C}}||p) &= \int_{\mathcal{C}} \frac{Z}{Z_{\mathcal{C}}} p(x) \log \frac{Z}{Z_{\mathcal{C}}} dx \\ 1050 \\ 1051 \quad &= \mathbb{P}_p(\mathcal{C}) \frac{Z}{Z_{\mathcal{C}}} \log \frac{Z}{Z_{\mathcal{C}}}, \end{aligned} \tag{B.16}$$

1052 where  $Z_{\mathcal{C}}$  satisfies  
 1053

$$\begin{aligned} 1054 \quad 1 &= \int_{\mathbb{R}^d} p_{\mathcal{C}} \\ 1055 \\ 1056 \quad &= \frac{Z}{Z_{\mathcal{C}}} \int_{\mathcal{C}} p(x) dx \\ 1057 \\ 1058 \quad &= \frac{Z}{Z_{\mathcal{C}}} \mathbb{P}_p(\mathcal{C}). \end{aligned} \tag{B.17}$$

1059 It follows that  
 1060

$$D(p_{\mathcal{C}}||p) = \log \frac{1}{\mathbb{P}_p(\mathcal{C})}. \tag{B.18}$$

1061 This value is found to be the minimizer of Problem (4.1) using Gibbs' inequality.  $\square$   
 1062

#### 1063 B.4 PROOF OF COROLLARY 1

1064 **Corollary** [Penalty methods] Penalty methods (2.5) cannot enforce  $\mathbb{P}_q(\mathcal{C}) = 1$ .  
 1065

1066 *Proof.* Penalty methods sample from  $p_\lambda$ , with finite  $\lambda$ . For all densities  $q \in \mathcal{P}_2(\mathbb{R}^d)$  satisfying the  
 1067 constraint  $\mathbb{P}_q(\mathcal{C}) = 1$ , the duality gap implies  
 1068

$$L(q_\lambda, \lambda) = g(\lambda) < F(p_{\mathcal{C}}) \leq F(q) = L(q, \lambda). \tag{B.19}$$

1069 Therefore,  $p_\lambda$  does not satisfy  $\mathbb{P}_q(\mathcal{C}) = 1$ .  $\square$   
 1070

1080 B.5 PROOF OF PROPOSITION 4  
10811082 **Proposition 7** [Variable splitting] Problem (4.1) is equivalent to the following problem:  
1083

1084 
$$\begin{aligned} & \underset{q \in \mathcal{P}_2(\mathbb{R}^d \times \mathcal{C})}{\text{minimize}} \quad D(q_x \| p) \\ & \text{subject to} \quad \mathbb{P}_q(x = z) = 1. \end{aligned} \tag{B.20}$$
  
1085

1086 *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  $x$ .  $\square$   
10871091 B.6 PROOF OF PROPOSITION 5  
10921093 **Proposition 8** [Attained duality] Strong duality holds and is attained for Problem (4.4).  
10941095 *Proof.* In order to apply Proposition 2.2 from Chamon et al. (2024), 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  $\mathcal{C}$ , making the non-negativity constraint of the density not qualified.  $\square$   
1096  
1097  
1098  
1099  
1100  
11011102 B.7 PROOF OF PROPOSITION 6  
11031104 Recall the relaxed problem  
1105

1106 
$$\begin{aligned} & \underset{q \in \mathcal{P}(\mathbb{R}^d \times \mathcal{C})}{\text{minimize}} \quad D(q \| p \otimes u_{\mathcal{C}}) + \rho \mathbb{E}_q [\|x - z\|^2] \\ & \text{subject to} \quad \mathbb{E}_q[x - z] = 0. \end{aligned} \tag{B.21}$$
  
1107  
1108

1109 **Proposition** [Problem approximation] The  $\rho$ -approximation converges to the strictly constrained  
1110 problem, as  
1111

1112 
$$q_{\rho} \xrightarrow[\rho \rightarrow +\infty]{\text{law}} p_{\mathcal{C}}.$$
  
1113

1114 *Proof of Proposition 6.* Recall that, because strong duality is attained, the solution of (4.4) is at-  
1115 tained by a distribution of the form

1116 
$$q_{\rho}(x, z) = \frac{1}{Z_{\lambda}} e^{-f(x)} e^{-\chi_{\mathcal{C}}(z)} e^{-\frac{\rho}{2} \|x - z\|^2} e^{-\lambda^T(x - z)} \tag{B.22}$$
  
1117  
1118

1119 Let  $z \in \mathcal{C}$  and  $x \neq z$  in  $\mathbb{R}^d$ . Then,  $q_{\rho}(x, z) \xrightarrow[\rho \rightarrow +\infty]{} 0 = p_{\mathcal{C}}^2(x, z)$ .  
11201121 Additionally,  
1122

1123 
$$q_{\rho}(z, z) = \frac{1}{Z_{\lambda}} e^{-f(z)} \tag{B.23}$$
  
1124  
1125  
1126  
1127  
1128  
1129  
1130  
1131  
1132  
1133

1134 C VARIATIONAL FRAMEWORK FOR LANGEVIN MONTE CARLO  
11351136 Consider the functional  
1137

1138 
$$F(q) = D(q\|p) = \int q \log(q/p). \quad (\text{C.1})$$

1139 The Wasserstein gradient flow is defined as the following differential system  
1140

1141 
$$\frac{\partial q}{\partial t} = \nabla \cdot \left( q \nabla \frac{\partial F}{\partial q} \right), \quad (\text{C.2})$$
  
1142

1143 For functional (C.1), the differential system is found to be  
1144

1145 
$$\frac{\partial q}{\partial t} = \nabla \cdot (q \nabla f(x)) + \Delta q(x, t), \quad (\text{C.3})$$
  
1146

1147 which is found to be the Fokker-Planck equation for the Langevin dynamics  
1148

1149 
$$dx = -\nabla f(x)dt + dB. \quad (\text{C.4})$$

1150 More details can be found in (Jordan et al., 1998; Ambrosio et al., 2008; Villani, 2021; Chamon  
1151 et al., 2024).  
1152  
1153  
1154  
1155  
1156  
1157  
1158  
1159  
1160  
1161  
1162  
1163  
1164  
1165  
1166  
1167  
1168  
1169  
1170  
1171  
1172  
1173  
1174  
1175  
1176  
1177  
1178  
1179  
1180  
1181  
1182  
1183  
1184  
1185  
1186  
1187

1188 **D CONNECTION BETWEEN GENERATIVE MODELS AND LANGEVIN**  
 1189 **SAMPLING**  
 1190

1191 **D.1 ENERGY-BASED MODELS (EBMs)**  
 1192

1193 An EBM defines a density

1194 
$$p(x) = \frac{1}{Z} \exp(-f_\theta(x)), \quad (\text{D.1})$$

1196 where  $f_\theta$  is a learned energy function. Sampling from  $p$  typically relies on Langevin dynamics (2.1)  
 1197 or stochastic gradient Langevin dynamics (SGLD) (Welling & Teh, 2011). EBMs with Langevin  
 1198 sampling have demonstrated strong performance across a range of tasks (Du & Mordatch, 2019),  
 1199 and offer distinct advantages over methods such as Variational Autoencoders (VAEs) (Kingma et al.,  
 1200 2013) and Generative Adversarial Networks (GANs) (Goodfellow et al., 2014). A particularly valuable  
 1201 property of EBMs is their flexibility in incorporating constraints via summing up the corresponding  
 1202 energies. From this perspective, our algorithm, when applied to EBMs, can be interpreted as providing  
 1203 stronger constraint enforcement through an augmented Lagrangian potential and corresponding  
 1204 proximal Langevin updates—going beyond the simple addition of constraint energies.  
 1205

1205 **D.2 SCORE-BASED GENERATIVE MODELS**  
 1206

1207 Score-based generative models aim to learn the score function  $\nabla \log p_t(x)$  of a family of progressively  
 1208 noised data distributions  $\{p_t\}_{t \in [0, T]}$ , rather than modeling the data density directly. Once the  
 1209 score is learned—typically via denoising score matching—samples can be generated by Langevin-  
 1210 type updates.

1211 **Annealed Langevin Dynamics** Proposed by Song & Ermon (2019), this method generates samples by applying Langevin dynamics at a sequence of decreasing noise levels  $\sigma_T > \dots > \sigma_1$ . 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  $\sigma_t$ . Then update step takes the form

1217 
$$x_{t+1} = x_t + \tau_t s_\theta(x, \sigma_t) + \sqrt{2\tau_t} w_t, \quad w_t \sim \mathcal{N}(0, I), \quad (\text{D.2})$$

1218 where  $\tau_t \propto \sigma_t^2$  are time-varying step sizes. The update takes the form of (A.7)  
 1219 with  $\nabla f(x, t) = -s_\theta(x, \sigma_t)$ . This can be seen as an unadjusted Langevin algorithm with temperature  
 1220  $\sigma_t$ , gradually refining the sample as noise decreases. In this case our algorithm can be  
 1221 directly applied at each noise level to impose constraints. It is worth noting that the projected diffusion  
 1222 model (Christopher et al., 2024) also falls into this category – a hard projection following each  
 1223 Langevin update within the annealed Langevin dynamics framework. Note that this covers the case  
 1224 where several Langevin steps are taken at fixed noise level, as in the work of Song & Ermon (2019),  
 1225 by choosing  $\tau_t$  to be constant for a number of steps  $t$ .  
 1226

1227 **D.3 DIFFUSION MODELS**

1228 **Denoising Diffusion Probabilistic Models (DDPM)** Denoising diffusion probabilistic models  
 1229 (DDPMs), introduced by Ho et al. (2020), define a forward process that gradually corrupts a  
 1230 data point  $y_0$  by adding Gaussian noise through a fixed Markov chain:

1232 
$$q(y_t | y_{t-1}) = \mathcal{N}(y_t; \sqrt{1 - \beta_t} y_{t-1}, \beta_t I), \quad (\text{D.3})$$

1233 where  $\beta_t \in (0, 1)$  is a small noise schedule. This leads to a closed-form expression for  $q(x_t | x_0)$ ,  
 1234 with the following definitions:

1235 
$$\alpha_t = 1 - \beta_t, \quad \bar{\alpha}_t = \prod_{s=1}^t \alpha_s. \quad (\text{D.4})$$

1238 The reverse process is parameterized by a neural network  $\epsilon_\theta(x_t, t)$ , which predicts the noise component.  
 1239 The sampling procedure follows:

1241 
$$x_{t+1} = \frac{1}{\sqrt{\alpha_t}} \left( x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t) \right) + \sigma_t w, \quad w_t \sim \mathcal{N}(0, I), \quad (\text{D.5})$$

1242 where  $\sigma_t$  is typically set to match the forward variance  $\beta_t$ . As noted by Ho et al. (2020), this  
 1243 step corresponds to an Euler-Maruyama discretization of a variant of Langevin dynamics, and the  
 1244 learned noise predictor  $\epsilon_\theta$  implicitly estimates the score  $\nabla \log p_t(x)$  up to a scaling factor. Hence,  
 1245 the sampling formula (D.5) takes the form (A.7) with  $\tau_t = \sigma_t^2/2$  and

$$1247 \quad \nabla \log p_t(x_t) \approx s_\theta(x_t, t) = -\frac{1}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t). \quad (\text{D.6})$$

1249 The DDPM can be regarded as a discrete score-based model under the variance preserving stochastic  
 1250 differential equation (VP-SDE) interpretation (Song et al., 2020), and thus our SAL sampling is valid  
 1251 for DDPM sampling.

#### 1253 D.4 SCORE-BASED DIFFUSION MODELS

1254 Score-based diffusion models (Song et al., 2020) directly learn the score function of perturbed data  
 1255 distributions and generate samples by simulating the reverse-time stochastic dynamics.

1257 **Forward SDE.** Define a forward Itô SDE that gradually adds noise to data  $x_0 \sim p_{\text{data}}$ :

$$1259 \quad dx = a(x, t) dt + b(t) dW_t, \quad (\text{D.7})$$

1260 where for the variance-preserving (VP) choice,

$$1262 \quad a(x, t) = -\frac{1}{2} \beta(t) x, \quad b(t) = \sqrt{\beta(t)}. \quad (\text{D.8})$$

1263 This yields marginal distributions  $p_t(x)$  that interpolate between the data and near-Gaussian noise  
 1264 as  $t$  increases.

1266 **Reverse SDE.** The time-reversed process follows

$$1268 \quad dx = [a(x, t) - b(t)^2 \nabla_x \log p_t(x)] dt + b(t) dW'_t, \quad (\text{D.9})$$

1269 where  $W'_t$  is a reverse-time Wiener process. A neural network  $s_\theta(x, t)$  is trained by score matching  
 1270 to approximate  $\nabla_x \log p(x, t)$ .

1271 **Predictor–Corrector sampling.** Once the score network is trained, our SAL sampling is applicable.  
 1272 SAL can also be integrated seamlessly into the predictor–corrector sampling scheme proposed by  
 1273 Song et al. (2020). The predictor–corrector sampler interleaves:

- 1275 • *Predictor*: an Euler–Maruyama step of the reverse SDE,

$$1277 \quad x_{t+1} = x_t - \tau_t [a(x_t, t) - b(t)^2 s_\theta(x_t, t)] + b(t) \sqrt{2\tau_t} w_t \quad w_t \sim \mathcal{N}(0, I_d). \quad (\text{D.10})$$

- 1278 • *Corrector*: a few steps of Langevin MCMC to refine samples,

$$1280 \quad x_{t+1} = x_t + \tau_t s_\theta(x_t, t) + \sqrt{2\tau_t} w_t, \quad w_t \sim \mathcal{N}(0, I_d). \quad (\text{D.11})$$

1281 Similar to the previous sections, these formulas take the form of (A.7), with different time-varying  
 1282 potential gradients  $\nabla f(x, t)$ .

1284 **Summary** Across EBMs, diffusion models, and hybrid schemes, the core sampling formula is an  
 1285 overdamped Langevin update, possibly annealed through noise scales. This makes our constrained  
 1286 sampling algorithm SAL compatible with all these approaches as a zero-shot plug-in.

1288  
 1289  
 1290  
 1291  
 1292  
 1293  
 1294  
 1295

1296 E EXPERIMENTAL DETAILS  
12971298 E.1 FIELD GENERATION  
12991300 **Baselines** Our algorithm is compared to projected Langevin Monte Carlo, primal-dual Monte  
1301 Carlo and constraint-penalized Langevin Monte Carlo. For the latter, we implement the variable-  
1302 splitting algorithm of Zhang et al. (2025), and the penalty parameter is a dual variable that is adapted  
1303 and updated following the same scheme as SAL.1304 **Sampling** Langevin Monte Carlo is iterated over 1000 steps, and we set  $\rho$  to follow a linear inter-  
1305 polation schedule between 2 and 20.1306 **Constraints** The field is subject to energy conservation (Example 2.1). The projection in closed  
1307 forms. For the penalty method, the penalty cost is  $c(x) = (\sum x_i - M)^2$ . The primal-dual Langevin  
1308 Monte Carlo algorithm enforces the constraint function  $h(x) = \sum x_i - M$  on average.1311 E.2 DATA ASSIMILATION  
13121313 **Context** In many geophysical and engineering applications, one relies on numerical simulation to  
1314 predict the time-dependent evolution of a complex system, whose state at physical time  $t$  is denoted  
1315 by  $x \in \mathbb{R}^d$ . But these models are inherently imperfect—either because of computational constraints  
1316 or incomplete knowledge of the true dynamics. When real-world observations  $y \in \mathbb{R}^m$  become  
1317 available (for example in digital-twin settings), we assume a statistical model of the form

1318 
$$y = h(x) + \varepsilon, \quad (\text{E.1})$$

1319 where  $h : \mathbb{R}^d \rightarrow \mathbb{R}^m$  is an observation operator and  $\varepsilon$  is the measurement error. The imperfect  
1320 simulation yields a prior forecast  $b \in \mathbb{R}^d$ , the background estimate, which must be adjusted using  $y$   
1321 to produce a more accurate estimate of the true state, usually referred to as the analysis, as the initial  
1322 condition for the next simulation. Equivalently, one seeks samples from the posterior

1323 
$$p(x | b, y) \propto p(y | x) p(x | b). \quad (\text{E.2})$$

1324 This estimation problem is formulated sequentially for each new observation, by propagating the  
1325 obtained posterior analysis with a forecast model, and repeating the process. Classically, this is  
1326 achieved by one of three approaches: sequential Monte Carlo methods (e.g. particle filters (Gor-  
1327 don et al., 1993)), ensemble-based filters (e.g. the Ensemble Kalman Filter (Evensen, 2003)), or  
1328 variational methods that solve for the MAP estimate (e.g. 3D-Var/4D-Var (Sasaki, 1970; Lorenc,  
1329 1986)). The 3D-Var algorithm assumes that the background error distribution and observation error  
1330 distribution are Gaussian,

1331 
$$x | b \sim \mathcal{N}(b, B), \quad \varepsilon \sim \mathcal{N}(0, P), \quad (\text{E.3})$$

1332 then taking negative logarithm of (E.2) yields the following optimization target:

1333 
$$J(x) = \frac{1}{2} \|y - h(x)\|_{P^{-1}}^2 + \frac{1}{2} \|x - b\|_{B^{-1}}^2. \quad (\text{E.4})$$

1334 Deep learning represents a promising tool to learn more complex priors for data assimilation (Huang  
1335 et al., 2024a; Rozet & Louppe, 2023; Qu et al., 2024; Blanke et al., 2024)1336 **Data** For simulating the Burgers equation, we implemented the same method as van Gastelen et al.  
1337 (2024), but we added an extra constant linear advection term. We work in Fourier space with the  
1338 first 20 Fourier modes. The field evolves according to the Burgers equation for 4 time units. We  
1339 generate 1,000 trajectories, with the field recorded at 10 timesteps for each trajectory, with the initial  
1340 state drawn at random in Fourier space with a power-law decay of the coefficient magnitude.1341 **Learning architecture** We implemented a DDPM diffusion model, using the formalism detailed  
1342 in Appendix D. Diffusion is learned in a latent space, defined as the first 10 Fourier modes. The  
1343 neural network involved is a fully connected network with depth 3 and width 128, using a cosine  
1344 time embedding. It is trained for 200 epochs. At sampling time, 1000 diffusion steps are used  
1345 with  $\rho = 10$ .

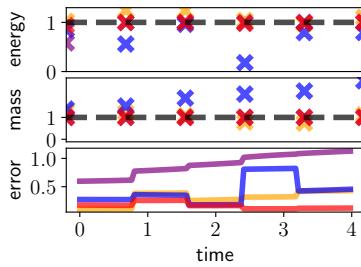
1350     **Baselines** Our algorithm is compared to unconstrained latent diffusion and Projected diffusion,  
 1351     which incorporate observations using diffusion posterior sampling Rout et al. (2023). We tried to  
 1352     apply penalty guidance diffusion in latent space but did not find a suitable method of incorporating  
 1353     observations in latent space.  
 1354

1355     **Constraints** The field is sampled subject to energy and mass conservation constraints (Exam-  
 1356     ple 2.1). The projection is computed by alternating projections on the two constraint set, which have  
 1357     closed forms.  
 1358

1359     The initial conditions are drawn at random following the same distribution of the training data.  
 1360

1361     **Additional results** Figure 6 shows the evolution of key metrics for a data assimilation trajectory,  
 1362     for the various methods compared.  
 1363

--- target    ■ Gaussian    ■ unconstrained    ■ projected    ■ SAL



1373     Figure 6: Mass conservation, energy conservation and  $\ell_2$  error.  
 1374

### 1376     E.3 FEASIBILITY PROBLEM 1377

1378     Trajectory planning for quadrotor obstacle avoidance is an imoprtant problem Le Hellard et al.  
 1379     (2025). We implement a linearized version of the planar quadrotor dynamics described in (Tedrake,  
 1380     2009).  
 1381

1382     **Data** The trajectories are discretized in time as  $(y_1, \dots, y_S, u_1, \dots, u_S) \in \mathbb{R}^{2S}$ , with  $S = 200$  and  
 1383     a time interval  $\Delta s = 0.01$ . The dynamics constraint  
 1384

$$1385 \quad \mathcal{C}_d := \{x \mid y(0) = 0, \dot{y}(s) = Ay(s) + Bu(s), |u(s)| \leq u_{\max}\} \quad (E.5)$$

1386     is described by a linear equality constraint, discretized into a linear system, and an inequality con-  
 1387     straint on the control inputs. The projection on this convex constraint set is obtained by Dykstra's  
 1388     double projection algorithm (Bílková & Šorel, 2021), and is used within the ADMM solver.  
 1389

1390     **Learning architecture** We implemented a DDPM diffusion model, using the formalism detailed  
 1391     in Appendix D. The trajectories signals are learned in a latent Fourier space encoding the first 10  
 1392     modes of the input signal. The neural network involved is a fully connected network with depth 3  
 1393     and width 128, using a cosine time embedding. It is trained for 200 epochs. At sampling time, 1000  
 1394     diffusion steps are used with  $\rho = 100$ .  
 1395

1396     **Baselines** We implement the latent projected diffusion algorithm Zampini et al. (2025), and diffu-  
 1397     sion guidance with constraint penalties (Carvalho et al., 2023), and propagate the penalty function  
 1398     through the decoder.  
 1399

1400     **Constraint** The obstacles are segments, and projecting onto the feasible region is performed by  
 1401     moving the penetrating trajectory portions trajectory either directly above or directly underneath  
 1402     the obstacle. For the penalty method, the constraint penalty is the quadratic distance between the  
 1403     trajectory and the obstacle, which is simple enough to differentiate through in this case.  
 1404

1404 E.4 ABLATION STUDIES FOR  $\rho$  AND  $\lambda$   
14051406 For the field generation task of Section 5.1, we generate 1000 samples of energy-constrained fields,  
1407 using different schedules for  $\rho$ . For each schedule, we evaluate the samples with the following  
1408 measure of error: we report the proportion of samples that fall near the unlikely positive mode of  
1409 the bimodal potential, implying that the sampled distribution deviates from the target1410 We experiment with 4 different schedules: two schedules use constant  $\rho$  throughout the iterations.  
1411 The two other schedules are linear and logarithmic interpolation between these two values. We run  
1412 the experiment for different numbers of Langevin iterations. The results are reported in the following  
1413 table.

1414 number of steps / schedule	1415 constant $\rho = \rho_{\min}$	1416 constant $\rho = \rho_{\max}$	1417 linear	1418 logarithmic
1415 1000	1416 42.4%	1417 22%	1418 0.04%	1419 3.8%
1416 5000	1417 44.7%	1418 2.8%	1419 0.001%	1420 0.0002%

1417 Table 2: Proportion of samples in the wrong mode for different schedules of  $\rho$ 1418 We observe that allowing  $\rho$  to vary across iterations substantially improves sample quality. With too  
1419 small  $\rho$ , the deviation between  $x$  and  $z$  is too large. With too large  $\rho$ , the chain fails to explore the  
1420 energy landscape. When the number of steps is limited, only annealed schedules manage to recover  
1421 the correct mode. This experiment highlights the importance of adaptive schedules in practice.1422 In the other experiments, we found that the time-varying step size induced by diffusion models, which  
1423 also scales  $\rho$ , was sufficient to balance exploration and constraint satisfaction1424 We conduct an ablation study on both the field generation experiment (Section 5.1) and the Burgers  
1425 data assimilation task (Section 5.2) to investigate the influence of the initial value of the dual  
1426 variables.1427 In the first experiment, the fields are sampled with the Langevin Monte Carlo algorithm, with fixed  
1428 potential. In the second experiment, a diffusion model is used, so the score function is time-varying  
1429 function throughout iterations.1430 We run our sampling algorithm with Gaussian initialization of  $\lambda$ , with different sizes  $\sigma$ . For each  
1431 value of  $\sigma$ , we run 100 independent chains. For the flow sampling experiment, we report the maxi-  
1432 mum number of sampling steps required to converge. For the data assimilation experiment, because  
1433 the dual problem changes over time, we do not evaluate the convergence of the dual variables. In-  
1434 stead, we report the average reconstruction accuracy.

1436 $\sigma$	1437 Steps to convergence	1438 Reconstruction error
1437 0	1438 10	1439 0.56
1438 1	1439 200	1440 0.56
1439 5	1440 800	1441 0.56
1440 10	1441 1000	1442 0.57
1441 20	1442 1600	1443 0.58
1442 50	1443 2500	1444 0.79
1443 100	1444 N/A	1445 1.34

1444 Table 3: Influence of  $\lambda$

1458     **Usage of Large Language Models** We used large language models at the sentence level to correct  
1459     English writing and avoid word repetition  
1460  
1461  
1462  
1463  
1464  
1465  
1466  
1467  
1468  
1469  
1470  
1471  
1472  
1473  
1474  
1475  
1476  
1477  
1478  
1479  
1480  
1481  
1482  
1483  
1484  
1485  
1486  
1487  
1488  
1489  
1490  
1491  
1492  
1493  
1494  
1495  
1496  
1497  
1498  
1499  
1500  
1501  
1502  
1503  
1504  
1505  
1506  
1507  
1508  
1509  
1510  
1511