

---

# Posterior Inference in Latent Space for Scalable Constrained Black-box Optimization

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Optimizing high-dimensional black-box functions under black-box constraints is  
2 a pervasive task in a wide range of scientific and engineering problems. These  
3 problems are typically harder than unconstrained problems due to hard-to-find  
4 feasible regions. While Bayesian optimization (BO) methods have been developed  
5 to solve such problems, they often struggle with the curse of dimensionality. Re-  
6 cently, generative model-based approaches have emerged as a promising alternative  
7 for constrained optimization. However, they suffer from poor scalability and are  
8 vulnerable to mode collapse, particularly when the target distribution is highly  
9 multi-modal. In this paper, we propose a new framework to overcome these chal-  
10 lenges. Our method iterates through two stages. First, we train flow-based models  
11 to capture the data distribution and surrogate models that predict both function  
12 values and constraint violations with uncertainty quantification. Second, we cast  
13 the candidate selection problem as a posterior inference problem to effectively  
14 search for promising candidates that have high objective values while not violating  
15 the constraints. During posterior inference, we find that the posterior distribution  
16 is highly multi-modal and has a large plateau due to constraints, especially when  
17 constraint feedback is given as binary indicators of feasibility. To mitigate this  
18 issue, we amortize the sampling from the posterior distribution in the latent space  
19 of flow-based models, which is much smoother than that in the data space. We  
20 empirically demonstrate that our method achieves superior performance on various  
21 synthetic and real-world constrained black-box optimization tasks. Our code is  
22 publicly available [here](#).

23 

## 1 Introduction

24 Optimizing high-dimensional black-box functions under black-box constraints is a fundamental task  
25 across numerous scientific and engineering problems, including machine learning [1], drug discovery  
26 [2, 3], control [4, 5], and industrial design [6, 7]. In most cases, these problems are much harder than  
27 unconstrained problems due to analytically undefined and hard-to-find feasible regions [8].

28 Bayesian Optimization (BO) has been widely used to solve black-box optimization problems in  
29 a sample-efficient manner [9, 10]. While most BO methods focus on unconstrained optimization  
30 problems, some works address problems with black-box constraints by developing new acquisition  
31 functions [1, 11] or relaxing the constraints [12, 13]. However, even without constraints, BO methods  
32 scale poorly to high dimensionality [14]. Moreover, incorporating constraints makes the function  
33 landscape highly complex, hindering accurate estimation of surrogate models.

34 Recently, generative models have emerged as an alternative solution for black-box optimization  
35 problems with constraints [15, 16, 17]. For example, we can leverage generative models to sample  
36 protein sequences that maximize the binding affinity while preserving the naturalness of the design.




Figure 1: Motivating figure. In a high-dimensional setting, sampling from the posterior distribution is beneficial for selecting candidates. However, the posterior distribution is highly multi-modal and has a large plateau due to the constraints (orange one). We can mitigate this issue by sampling latents from the posterior distribution (purple one) of the latent space and projecting them into the data space.

37 However, existing methods rely on MCMC-based approaches [15], which limit scalability in high-  
 38 dimensional spaces. While one can fine-tune pretrained generative models with reward functions  
 39 [18, 19, 20], naive application of fine-tuning methods is vulnerable to mode collapse when the target  
 40 distribution is highly multi-modal [21], which leads to a convergence on sub-optimal solutions.

41 In this paper, we propose a novel generative model-based framework for constrained black-box  
 42 optimization to overcome the aforementioned limitations. To efficiently explore high-dimensional  
 43 spaces, we first frame the candidate selection problem as sampling from the posterior distribution,  
 44 which can be constructed by multiplying the prior distribution with a Lagrangian-relaxed objective.  
 45 To sample candidates from the posterior distribution, our key idea is to amortize inference in the latent  
 46 space of a flow-based model using an outsourced diffusion sampler [21], as illustrated in Figure 1.  
 47 Since the posterior distribution in the latent space is much smoother than that in the data space, we  
 48 can approximate the distribution more accurately and alleviate the mode collapse problem [22].

49 Our method iterates through two stages. First, we train a flow-based model to capture the current  
 50 data distribution and surrogate models to predict the objective value and constraints, respectively.  
 51 For the surrogate models, we use an ensemble to quantify the uncertainty of the prediction, as we  
 52 have only a small amount of data that covers a tiny fraction of the whole search space. We treat a  
 53 trained flow-based model as a prior, and Lagrangian relaxation of the objective as a reward function.  
 54 Second, we sample candidates from the posterior distribution. As the posterior distribution is highly  
 55 multi-modal and has a large plateau due to constraints, especially when constraint feedback is given  
 56 as binary indicators of feasibility, we train a diffusion sampler that amortizes the posterior distribution  
 57 in the latent space of flow models. Then, we sample latents from the diffusion sampler and project  
 58 them into data space using a deterministic mapping derived from the trained flow model. By repeating  
 59 these two stages, we can effectively capture high-scoring regions that satisfy the constraints.

60 We conduct extensive experiments on three synthetic and three real-world benchmarks to validate the  
 61 superiority of our method on scalable constrained black-box optimization problems. We also consider  
 62 a more challenging scenario where the feedback from the constraints is given as a binary value. We  
 63 empirically show that our method outperforms several competitive baselines across different tasks.

## 64 2 Related Works

### 65 2.1 Constrained Black-box Optimization

66 Most scientific and engineering optimization problems involve black-box constraints, such as the  
 67 synthesizability of molecules in chemical design [2] and safety constraints in robot control policies  
 68 [4]. Existing BO methods solve this problem by either integrating the constraints directly into the  
 69 acquisition function (cEI [23], LogcEI [24]) or by employing trust region approaches for scalability  
 70 (SCBO [8], PCAGP-SCBO [7]). Another line of work utilizes evolutionary algorithms like CMA-  
 71 ES [25, 26] with an augmented Lagrangian method to navigate constrained spaces. However, the  
 72 performance of these methods often degrades as dimensionality and the number of evaluations  
 73 increase, which motivates the need for a more scalable approach.

74 **2.2 Generative Model-based Optimization**

75 There are several attempts to utilize generative models for black-box optimization. In an offline  
 76 setting, DDOM [27] trains a conditional diffusion model with classifier-free guidance and applies a  
 77 loss-reweighting to emphasize samples with high objective values. DiffOPT [15] solves a constrained  
 78 optimization problem. It applies diffusion to capture data distribution, followed by an iterative  
 79 importance-sampling procedure. In an online setting, DiffBBO [28] and DiBO [29] both leverage  
 80 diffusion models and incorporate uncertainty estimation during candidate selection. DiBO treats  
 81 candidate selection as posterior inference to guide sampling toward regions of high reward and uncer-  
 82 tainty, while DiffBBO selects conditioning targets by employing an uncertainty-based acquisition  
 83 function. Unfortunately, constrained black-box optimization in the online setting remains unexplored.

84 **2.3 Amortized Inference in Flow-based and Diffusion Models**

85 Given a diffusion or flow prior  $p_\theta(\mathbf{x})$  trained on a dataset and a reward function  $r(\mathbf{x})$ , sampling from  
 86 the posterior  $p_{\text{post}}(\mathbf{x}) \propto p_\theta(\mathbf{x})r(\mathbf{x})$  has numerous applications in downstream tasks [30, 18, 31, 32,  
 87 19, 21]. However, direct sampling from the unnormalized posterior  $p_\theta(\mathbf{x})r(\mathbf{x})$  is intractable [18, 33].

88 To address this problem, some approaches train classifiers directly within intermediate noised spaces  
 89 [30, 34] while others approximate posterior sampling via Markov Chain Monte Carlo (MCMC)  
 90 procedures [22, 32, 35, 36]. However, training classifiers in noisy data spaces and employing MCMC  
 91 methods scale poorly to high dimensionality. On the other hand, several methods utilize reinforcement  
 92 learning [37, 38] or stochastic optimal control [19] to fine-tune the pretrained model and amortize the  
 93 posterior sampling. Meanwhile, naive implementations of fine-tuning methods can be prone to mode  
 94 collapse when the target distribution is highly multi-modal and has a large plateau region [21].

95 To mitigate this issue, we adopt the outsourced diffusion sampler method proposed by Venkatraman  
 96 et al. [21]. Matching the distribution within the latent space significantly simplifies the alignment  
 97 task when the distribution is highly multi-modal and has a large flat region in the original data space.

98 **3 Preliminaries**

99 **3.1 Constrained Black-box Optimization**

100 In constrained black-box optimization, our problem is:

$$\text{find } \mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \quad \text{s.t. } g^{(1)}(\mathbf{x}) \leq 0, \dots, g^{(M)}(\mathbf{x}) \leq 0$$

with  $R$  rounds of  $B$  batch of queries (1)

101 The objective function  $f : \mathcal{X} \rightarrow \mathbb{R}$  and constraints  $g^{(1)}, \dots, g^{(M)} : \mathcal{X} \rightarrow \mathbb{R}$  are black-box functions.  
 102 We also consider a more challenging scenario, only access to information on whether we violate  
 103 constraints or not, i.e.,  $h^{(m)}(\mathbf{x}) = \mathbb{I}[g^{(m)}(\mathbf{x}) > 0]$ . We refer to this as an indicator constraint.

104 **3.2 Flow-based Models**

105 Flow-based models [39, 40, 41] are a class of generative models for approximating a target distribution  
 106  $q(\mathbf{x})$ . Flow-based models are defined via the deterministic ordinary differential equation (ODE):

$$d\mathbf{x}_t = v_\theta(\mathbf{x}_t, t) dt \quad (2)$$

107 where  $v_\theta(\mathbf{x}_t, t) : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$  is a parametric velocity field.

108 For each given velocity field, the corresponding flow  $\psi_\theta(\mathbf{x}_0, t) : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$  satisfies:

$$\frac{d}{dt} \psi_\theta(\mathbf{x}_0, t) = v_\theta(\psi_\theta(\mathbf{x}_0, t), t), \quad \psi_\theta(\mathbf{x}_0, 0) = \mathbf{x}_0. \quad (3)$$

109 The velocity field  $v_\theta(\mathbf{x}_t, t)$  defines a continuous probability path  $p_t$  induced by the flow:

$$\mathbf{x}_t = \psi_\theta(\mathbf{x}_0, t) \sim p_t, \quad \text{where } \mathbf{x}_0 \sim p_0. \quad (4)$$

110 **Training Flow-based Models.** We use Flow Matching [39] to learn the velocity field  $v_\theta$  that generates  
 111 a path interpolating smoothly between an initial distribution  $p_0 = p$  and a target distribution  $p_1 = q$ .




Figure 2: Overview of our method. **Phase 1:** Train flow-based models and proxies for the objective and constraints. **Phase 2:** Sample candidates from the posterior distribution using an outsourced diffusion sampler. After sampling, we utilize filtering to enhance sample efficiency. Then, we evaluate samples, update the dataset, and repeat the process until the evaluation budget is exhausted.

112 We employ the simplest linear interpolation path  $\mathbf{x}_t = (1-t)\mathbf{x}_0 + t\mathbf{x}_1$ , with derivative  $\frac{d\mathbf{x}_t}{dt} = \mathbf{x}_1 - \mathbf{x}_0$ ,  
 113 following [39]. The Flow Matching loss is expressed as:

$$\mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}_{\mathbf{x}_0 \sim \mathcal{N}(0, I), \mathbf{x}_1 \sim q(\mathbf{x}), t \sim \text{Unif}(0, 1)} [\|\mathbf{v}_\theta(\mathbf{x}_t, t) - (\mathbf{x}_1 - \mathbf{x}_0)\|_2^2]. \quad (5)$$

### 114 3.3 Posterior Inference in Flow-based and Diffusion Models

115 Given a pretrained flow-based prior  $p_\theta(\mathbf{x})$ , and a reward function  $r(\mathbf{x})$ , we consistently encounter a  
 116 situation where we need to sample from the posterior distribution,  $p_{\text{post}}(\mathbf{x}) \propto p_\theta(\mathbf{x})r(\mathbf{x})$ . To sample  
 117 from this intractable [33] distribution, we utilize the outsourced diffusion sampling [21].

118 We can interpret the sampling process of flow-based models into a noise generation  $\mathbf{z} \sim p(\mathbf{z})$ , followed  
 119 by a deterministic transformation  $\mathbf{x} = f_\theta(\mathbf{z})$ , where  $p(\mathbf{z})$  is standard normal and  $f_\theta$  represents the  
 120 learned mapping derived by prior. Under this formulation, by Proposition 3.1 of [21], we can sample  
 121 from the posterior distribution by substituting noise generation as  $\mathbf{z} \sim p_{\text{post}}(\mathbf{z}) \propto p(\mathbf{z})r(f_\theta(\mathbf{z}))$ .

122 To approximate the target distribution  $p_\psi(\mathbf{z}) \approx p_{\text{post}}(\mathbf{z})$ , we can learn the parameters of diffusion  
 123 sampler  $\psi$  with the trajectory balance (TB) objective [42, 43]:

$$\mathcal{L}_{\text{TB}}(\mathbf{z}_{0:1}; \psi) = \left( \log \frac{Z_\psi p(\mathbf{z}_0) \prod_{i=0}^{T-1} p_F(\mathbf{z}_{(i+1)\Delta t} | \mathbf{z}_{i\Delta t}; \psi)}{p(\mathbf{z}_1) r(f_\theta(\mathbf{z}_1)) \prod_{i=1}^T p_B(\mathbf{z}_{(i-1)\Delta t} | \mathbf{z}_{i\Delta t})} \right)^2, \quad (6)$$

124 where  $Z_\psi$  is the parameterized partition estimator,  $(\mathbf{z}_0 \rightarrow \mathbf{z}_{\Delta t} \rightarrow \dots \mathbf{z}_1 = \mathbf{z})$  is the discrete  
 125 time Markov chain of reverse-time stochastic differential equation (SDE) [44] with time increment  
 126  $\Delta t = \frac{1}{T}$ .  $p_F$  and  $p_B$  are transition kernels of the discretized reverse and forward SDE.

## 127 4 Method

128 In this section, we introduce **CiBO**, a new framework for scalable constrained black-box optimization  
 129 by leveraging generative models. Our method consists of two iterative stages. First, we train a  
 130 flow-based model to capture the data distribution and surrogate models to predict objective values  
 131 and constraints with uncertainty quantification. Next, we sample candidates from the posterior  
 132 distribution. To accomplish this, we train a diffusion sampler that draws samples from the posterior  
 133 distribution in the latent space. After sampling, we evaluate candidates, update the dataset, and repeat  
 134 the process until the evaluation budget is exhausted. Figure 2 illustrates the overview of our method.

### 135 4.1 Phase 1. Training Models

136 In each round  $r$ , we have a pre-collected dataset  $\mathcal{D}_r = \{\mathbf{x}_i, y_i, \mathbf{c}_i\}_{i=1}^I$ , where  $y_i = f(\mathbf{x}_i)$ ,  $\mathbf{c}_i =$   
 137  $\{c_i^m | c_i^m = g^{(m)}(\mathbf{x}_i), \forall m = 1, \dots, M\}$ , and  $I$  is the number of data points collected so far.

138 **Training Prior.** We first train a prior model  $p_\theta$  to capture the current data distribution. As the search  
 139 space is too high-dimensional, it is better to implicitly constrain the search space close to the current  
 140 data distribution. We use flow-based models to learn this distribution using Equation (5).

141 **Training Surrogates.** We also train surrogate models to predict both objective values and constraints.  
 142 As we are only able to access a small number of data points in the vast search space, we need to

143 properly quantify the uncertainty of the prediction. To this end, we train an ensemble of proxies to  
 144 estimate objective values with uncertainty quantification [45]. Specifically, we train an ensemble of  
 145  $K$  proxies  $f_{\phi_1}, \dots, f_{\phi_K}$  for objective values, and individual proxy  $g_{\phi}^{(1)}, \dots, g_{\phi}^{(M)}$  for each constraint.

146 **Reweighted Training.** During training, we introduce a reweighted training scheme [27, 46, 47] to  
 147 focus on promising data points with high objective values while not violating constraints. Specifically,  
 148 the weight for each data point is computed as follows:

$$l(y, \mathbf{c}) = y - \lambda \sum_{m=1}^M \max(0, c^m), \quad w(y, \mathbf{c}, \mathcal{D}_r) = \frac{\exp(l(y, \mathbf{c}))}{\sum_{(y', \mathbf{c}') \in \mathcal{D}_r} \exp(l(y', \mathbf{c}'))}. \quad (7)$$

149 Then, our training objective for flow-based models and proxies can be described as follows:

$$\mathcal{L}(\theta) = \mathbb{E}_{\mathbf{x}_0 \sim \mathcal{N}(0, I), (\mathbf{x}, y, \mathbf{c}) \in \mathcal{D}_r, t \sim \text{Unif}(0, 1)} [w(y, \mathbf{c}, \mathcal{D}_r) \|v_{\theta}(\mathbf{x}_t, t) - (\mathbf{x} - \mathbf{x}_0)\|_2^2], \quad (8)$$

$$\mathcal{L}(\phi) = \sum_{(\mathbf{x}, y, \mathbf{c}) \in \mathcal{D}_r} w(y, \mathbf{c}, \mathcal{D}_r) \left[ \sum_{k=1}^K (y - f_{\phi_k}(\mathbf{x}))^2 + \sum_{m=1}^M (c^m - g_{\phi}^{(m)}(\mathbf{x}))^2 \right]. \quad (9)$$

## 151 4.2 Phase 2. Sampling Candidates

152 After training models, we proceed to select candidates for evaluation in the current round. As the  
 153 search space is high-dimensional, the prediction of surrogate models is likely to be inaccurate in  
 154 regions that are too far away from the dataset collected so far. Therefore, it is advantageous to sample  
 155 candidates from the distribution that satisfies the two desiderata: (1) promote exploration towards  
 156 high-scoring and feasible regions, and (2) prevent sampling candidates that deviate too far from the  
 157 current data distribution. To accomplish these objectives, we cast the candidate selection problem as  
 158 sampling from the target distribution  $p_{\text{post}}$  defined as follows:

$$p_{\text{post}}(\mathbf{x}) = \arg \max_{p \in \mathcal{P}} \mathbb{E}_{\mathbf{x} \sim p} [r_{\phi}(\mathbf{x})] - \frac{1}{\beta} \cdot D_{\text{KL}}(p \| p_{\theta}), \quad (10)$$

159 where  $\mathcal{P}$  is the space of all probability distributions over the domain  $\mathcal{X}$ , and

$$r_{\phi}(\mathbf{x}) = \mu_{\phi}(\mathbf{x}) + \gamma \cdot \sigma_{\phi}(\mathbf{x}) - \lambda \sum_{m=1}^M \max(0, g_{\phi}^{(m)}(\mathbf{x})). \quad (11)$$

160  $\mu_{\phi}(\mathbf{x})$  and  $\sigma_{\phi}(\mathbf{x})$  represent the mean and standard deviation from the ensemble of surrogate models  
 161 for the objective.  $\gamma$  controls exploration-exploitation trade-off,  $\beta$  is an inverse temperature, and  $\lambda$  is a  
 162 Lagrange multiplier. Based on derivation from [48], our target distribution analytically derived as:

$$p_{\text{post}}(\mathbf{x}) \propto p_{\theta}(\mathbf{x}) \exp(\beta \cdot [r_{\phi}(\mathbf{x})]). \quad (12)$$

163 If we treat the flow-based model  $p_{\theta}(\mathbf{x})$  as a prior and the exponential term  $\exp(\beta \cdot [r_{\phi}(\mathbf{x})])$  as a  
 164 reward  $r(\mathbf{x})$ , then our objective is to sample from the posterior distribution  $p_{\text{post}}(\mathbf{x}) \propto p_{\theta}(\mathbf{x})r(\mathbf{x})$ .

165 **Amortized Inference in Latent Space.** However, directly sampling from this posterior is in-  
 166 tractable [33]. Also, the target posterior is highly multi-modal and has a large plateau due to the  
 167 constraint penalties in (11), making finetuning-based methods [18, 37] susceptible to mode collapse.  
 168 To this end, we utilize an amortized sampler in the latent space suggested by Venkatraman et al [21].

169 As introduced in Section 3.3, we can view the sampling procedure of flow-based models as drawing  
 170 samples from the standard normal distribution  $\mathbf{z} \sim p(\mathbf{z})$ , followed by the deterministic transformation  
 171  $\mathbf{x} = f_{\theta}(\mathbf{z})$ . Within this framework, we can generate samples from the posterior distribution  $p_{\text{post}}(\mathbf{x})$   
 172 by modifying the noise generation distribution as follows:

$$\mathbf{z} \sim p_{\text{post}}(\mathbf{z}) \propto p(\mathbf{z})r(f_{\theta}(\mathbf{z})). \quad (13)$$

173 To sample latents  $\mathbf{z}$  from the posterior distribution in the latent space  $p_{\text{post}}(\mathbf{z})$ , we train a diffusion  
 174 model  $p_{\psi}(\mathbf{z})$  to amortize  $p_{\text{post}}(\mathbf{z})$  with the following Trajectory Balance (TB) objective:

$$\mathcal{L}_{\text{TB}}(\mathbf{z}_{0:1}; \psi) = \left( \log \frac{Z_{\psi} p(\mathbf{z}_0) \prod_{i=0}^{T-1} p_F(\mathbf{z}_{(i+1)\Delta t} | \mathbf{z}_{i\Delta t}; \psi)}{p(\mathbf{z}_1) r(f_{\theta}(\mathbf{z}_1)) \prod_{i=1}^T p_B(\mathbf{z}_{(i-1)\Delta t} | \mathbf{z}_{i\Delta t})} \right)^2. \quad (14)$$

175 By training an amortized sampler in the latent space of flow-based models, we can more accurately  
 176 sample candidates from the target distribution as the posterior distribution in the latent space is  
 177 smoother than that in the data space. We also adopt off-policy training, detailed in Appendix D.2.1




Figure 3: Comparison between our method and baselines in synthetic tasks. Experiments are conducted with four random seeds, and the mean and one standard deviation are reported. A dashed line means that no feasible solutions were found.

### 178 4.3 Filtering, Evaluation and Moving Dataset

179 **Filtering.** After sampling from the posterior distribution, we need to carefully select candidates for  
180 the sample efficiency of the algorithm. To do so, we generate  $N \cdot B$  samples from the amortized  
181 sampler and select the top- $B$  samples in terms of Lagrangian relaxation of objectives as candidates.

182 **Evaluation and Moving Dataset.** We evaluate the values of the objective function and constraint  
183 functions for each selected candidate, then update the dataset with new observations. During the  
184 update, we empirically find that taking only a subset of total observations is beneficial in terms of  
185 computational complexity. We remove the samples with the lowest Lagrangian-relaxed objective if  
186 the dataset size is larger than the buffer size  $L$ . The pseudocode of our method is in Algorithm 1.

## 187 5 Experiments

188 In this section, we report experimental results for scalable constrained black-box optimization tasks.  
189 First, we perform experiments on three 200-dimensional synthetic functions, which are the standard  
190 benchmarks in Bayesian Optimization (BO) studies [14]. Furthermore, we assess the performance of  
191 our method on a more challenging scenario, where the feedback from constraints is given as binary  
192 indicators of feasibility. We refer to this setting as the indicator constraint setting. Finally, we conduct  
193 experiments on three real-world optimization tasks: Rover Planning 60D [8, 49], Mopta 124D [6],  
194 and Lasso DNA 180D [50]. The detailed description of each task can be found in Appendix A.

195 For evaluation, we report the minimum regret of feasible solutions over the course of the training,  
196 and assign the largest regret found in all algorithms to the infeasible solutions, following [8, 51].

### 197 5.1 Baselines

198 We compare our method with several constrained BO baselines, including cEI [23], LogcEI [24],  
199 SCBO [8], PCAGP-SCBO [7], and the evolutionary search algorithm CMA-ES [25]. We also evaluate  
200 generative model-based approaches specifically designed for constrained optimization: DiffOPT [15],  
201 as well as methods that can be extended to constrained optimization via the Lagrangian relaxation:  
202 DDOM [27] and DIBO [29]. Detailed implementations of all baselines are provided in Appendix B.




Figure 4: Comparison between our method and baselines in real-world tasks. Experiments are conducted with four random seeds, and the mean and one standard deviation are reported. A dashed line means that no feasible solutions were found.

## 203 5.2 Synthetic Experiments

204 We first conduct experiments on three synthetic functions, Rastrigin-200D, Ackley-200D, and  
 205 Rosenbrock-200D. For each function, we utilize two inequality constraints proposed by SCBO [8]:  
 206  $\sum_{d=1}^{200} x_d \leq 0$  and  $\|\mathbf{x}\|_2^2 \leq 30$ . We conduct all experiments with an initial dataset size of  $|D_0| = 200$ ,  
 207 using a batch size of  $B = 100$  and a maximum evaluation limit of 10,000. In the indicator constraints  
 208 scenarios, as it is too challenging to find an initial feasible solution across all baselines, we sample 10  
 209 points within feasible regions during initialization.

210 As shown in the Figure 3, our method outperforms all baselines across different synthetic tasks, both  
 211 in the standard and indicator constraints. Generative model-based methods, including DiffOPT and  
 212 DDOM, struggle to find a feasible solution and fail to improve on indicator constraints. While DiBO  
 213 achieves better feasibility, its finetuning-based approach suffers from mode collapse and tends to  
 214 converge to suboptimal solutions. These results show that employing an outsourced diffusion sampler  
 215 significantly enhances performance in constrained black-box optimization by effectively capturing  
 216 multi-modal and expansive flat target distributions.

217 Constrained BO methods (SCBO, PCAGP-SCBO, and LogcEI) successfully identify feasible points  
 218 but show limited sample efficiency compared to our method across all tasks. The evolutionary search  
 219 algorithm CMA-ES performs modestly in general but fails to find a feasible solution for some tasks.  
 220 These results underscore that our approach effectively captures both high-scoring and feasible regions  
 221 in a sample-efficient manner. Furthermore, compared to other baselines, our method consistently  
 222 finds feasible solutions throughout the optimization process, which is illustrated in Appendix E.1.

## 223 5.3 Real World Experiments

224 To validate the robustness of our approach, we evaluate our method on three challenging real-world  
 225 benchmark problems: (1) Rover Planning in 60 dimensions with 15 infeasible square-shaped regions,  
 226 (2) Mopta in 124 dimensions with 68 constraints, and (3) Lasso DNA in 180 dimensions with 5  
 227 constraints. For all experiments, we initialize with  $|D_0| = 200$  data points and limit evaluations to  
 228 2,000. We use a batch size of  $B = 50$  for Rover Planning and Lasso DNA, and  $B = 20$  for Mopta,  
 229 as no baseline methods could identify feasible solutions with  $B = 50$ .

230 As illustrated in Figure 4, our approach consistently identifies high-quality feasible solutions with  
 231 superior sample efficiency across all tasks. We observe that the performance gap between our method  
 232 and other baselines becomes larger on real-world problems and most baselines failed to find any  
 233 feasible solutions for the challenging Mopta-124D and DNA-180D tasks. While SCBO is the only  
 234 competing method to achieve feasibility alongside our approach, it exhibits lower sample efficiency.  
 235 This highlights the robustness of our approach for scalable constrained black-box optimization.

## 236 5.4 Additional Analysis

237 In this section, we conduct a comprehensive analysis of each component of our proposed method  
 238 through ablation experiments on Rastrigin 200D and Rover Planning 60D.




Figure 5: Additional analysis for various components of CiBO. Experiments are conducted with four random seeds, and the mean and one standard deviation are reported.

239 **Reweighted Training.** To investigate the effectiveness of our reweighted training approach suggested  
 240 in Equation (7), we conduct a comparative analysis of two variants: training without reweighting,  
 241 applying weights based on the objective values (Objective-prioritized). As shown in Figure 5a,  
 242 variants without reweighting or using objective-prioritized reweighting exhibit low sample efficiency.

243 **Sampling Procedure.** We analyze the effect of each component in candidate sampling. We conduct  
 244 experiments with two variants: removing filtering, and removing both filtering and the diffusion  
 245 sampler, thus sampling candidates directly from the prior  $p_\theta$ . As depicted in Figure 5b, there is a  
 246 significant performance gap between our method and other variants, validating the effectiveness of  
 247 each proposed component. We also experiment with the filtering coefficient  $N$  in Appendix E.2.

248 **Lagrangian Multiplier  $\lambda$ .** We introduce the Lagrangian multiplier  $\lambda$ . As shown in Figure 5c, setting  
 249  $\lambda = 0$  (eliminating the constraint penalty) significantly degrades performance on both tasks, as it only  
 250 focuses on high objective values and neglects the feasibility of solutions. Conversely, excessively  
 251 high  $\lambda$  values diminish the influence of the objective function, resulting in reduced sample efficiency.

252 **Inverse Temperature  $\beta$ .** The inverse temperature controls the balance between the prior  $p_\theta(\mathbf{x})$  and  
 253 the reward function  $r(\mathbf{x})$ . We conduct experiments by varying  $\beta$  values. As shown in Figure 5d, using  
 254 a moderately high  $\beta$  generally helps to improve sample efficiency. However, if  $\beta$  is too high, the  
 255 performance is heavily dependent on the accuracy of surrogate models, leading to slow convergence.  
 256 This validates that incorporating prior distribution is crucial for scalability (Section 4.2).

257 **Further Analysis.** To further understand our method, we analyze the impact of the buffer size  $L$   
 258 (Appendix E.3), batch size  $B$ , and initial dataset size  $|\mathcal{D}_0|$  (Appendix E.5). We also investigate the  
 259 effect of off-policy training (Appendix E.4) and runtime scalability of our method (Appendix F).

## 260 6 Conclusion

261 We introduced CiBO, a generative model-based framework for scalable constrained black-box  
 262 optimization. Our approach formulates candidate selection as posterior inference, leveraging flow-  
 263 based models to capture the data distribution and surrogate models to predict both objectives and  
 264 constraints. By amortizing posterior sampling in the latent space with outsourced diffusion samplers,  
 265 our method effectively addresses the challenges posed by highly multi-modal and flat posterior  
 266 distributions that arise from incorporating constraints. Extensive experiments across synthetic and  
 267 real-world benchmarks demonstrate the superiority of our proposed method.

268 **References**

269 [1] Jacob Gardner, Matt Kusner, Kilian Weinberger, John Cunningham, et al. Bayesian optimization  
270 with inequality constraints. In *International Conference on Machine Learning*, pages 937–945.  
271 PMLR, 2014.

272 [2] Ryan-Rhys Griffiths and José Miguel Hernández-Lobato. Constrained bayesian optimization for  
273 automatic chemical design using variational autoencoders. *Chemical science*, 11(2):577–586,  
274 2020.

275 [3] Ksenia Korovina, Sailun Xu, Kirthevasan Kandasamy, Willie Neiswanger, Barnabas Poczos,  
276 Jeff Schneider, and Eric Xing. Chembo: Bayesian optimization of small organic molecules  
277 with synthesizable recommendations. In *International Conference on Artificial Intelligence and  
278 Statistics*, pages 3393–3403. PMLR, 2020.

279 [4] Felix Berkenkamp, Angela P Schoellig, and Andreas Krause. Safe controller optimization for  
280 quadrotors with gaussian processes. In *International conference on robotics and automation  
281 (ICRA)*, 2016.

282 [5] Felix Berkenkamp, Andreas Krause, and Angela P Schoellig. Bayesian optimization with safety  
283 constraints: safe and automatic parameter tuning in robotics. *Machine Learning*, 112(10):3713–  
284 3747, 2023.

285 [6] MF Anjos and DR Jones. Mopta 2008 benchmark. *URL* <http://www.miguelanjos.com/jones-benchmark>, 2009.

287 [7] Hauke F Maathuis, Roeland De Breuker, and Saullo GP Castro. High-dimensional bayesian  
288 optimisation with large-scale constraints via latent space gaussian processes. *arXiv preprint  
289 arXiv:2412.15679*, 2024.

290 [8] David Eriksson and Matthias Poloczek. Scalable constrained bayesian optimization. In *International  
291 conference on artificial intelligence and statistics*, pages 730–738. PMLR, 2021.

292 [9] Peter I Frazier. A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.

293 [10] Roman Garnett. *Bayesian optimization*. Cambridge University Press, 2023.

294 [11] José Miguel Hernández-Lobato, Michael Gelbart, Matthew Hoffman, Ryan Adams, and Zoubin  
295 Ghahramani. Predictive entropy search for bayesian optimization with unknown constraints. In  
296 *International conference on machine learning*, pages 1699–1707. PMLR, 2015.

297 [12] Victor Picheny, Robert B Gramacy, Stefan Wild, and Sébastien Le Digabel. Bayesian optimization  
298 under mixed constraints with a slack-variable augmented lagrangian. *Advances in neural  
299 information processing systems*, 29, 2016.

300 [13] Setareh Ariaifar, Jaume Coll-Font, Dana Brooks, and Jennifer Dy. Admmbo: Bayesian op-  
301 timization with unknown constraints using admm. *Journal of Machine Learning Research*,  
302 20(123):1–26, 2019.

303 [14] David Eriksson, Michael Pearce, Jacob Gardner, Ryan D Turner, and Matthias Poloczek.  
304 Scalable global optimization via local bayesian optimization. In *Advances in Neural Information  
305 Processing Systems (NeurIPS)*, 2019.

306 [15] Lingkai Kong, Yuanqi Du, Wenhao Mu, Kirill Neklyudov, Valentin De Bortoli, Dongxia Wu,  
307 Haorui Wang, Aaron M Ferber, Yian Ma, Carla P Gomes, and Chao Zhang. Diffusion models  
308 as constrained samplers for optimization with unknown constraints. In *The 28th International  
309 Conference on Artificial Intelligence and Statistics*, 2025.

310 [16] Wenqian Xing, JungHo Lee, Chong Liu, and Shixiang Zhu. Black-box optimization with  
311 implicit constraints for public policy. In *Proceedings of the AAAI Conference on Artificial  
312 Intelligence*, volume 39, pages 28511–28519, 2025.

313 [17] Masatoshi Uehara, Xingyu Su, Yulai Zhao, Xiner Li, Aviv Regev, Shuiwang Ji, Sergey Levine,  
314 and Tommaso Biancalani. Reward-guided iterative refinement in diffusion models at test-time  
315 with applications to protein and dna design. *arXiv preprint arXiv:2502.14944*, 2025.

316 [18] Siddarth Venkatraman, Moksh Jain, Luca Scimeca, Minsu Kim, Marcin Sendera, Mohsin  
 317 Hasan, Luke Rowe, Sarthak Mittal, Pablo Lemos, Emmanuel Bengio, Alexandre Adam, Jarrid  
 318 Rector-Brooks, Yoshua Bengio, Glen Berseth, and Nikolay Malkin. Amortizing intractable  
 319 inference in diffusion models for vision, language, and control. In *The Thirty-eighth Annual  
 320 Conference on Neural Information Processing Systems*, 2024.

321 [19] Carles Domingo-Enrich, Michal Drozdzal, Brian Karrer, and Ricky T. Q. Chen. Adjoint  
 322 matching: Fine-tuning flow and diffusion generative models with memoryless stochastic optimal  
 323 control. In *The Thirteenth International Conference on Learning Representations*, 2025.

324 [20] Masatoshi Uehara, Yulai Zhao, Kevin Black, Ehsan Hajiramezanali, Gabriele Scalia,  
 325 Nathaniel Lee Diamant, Alex M Tseng, Tommaso Biancalani, and Sergey Levine. Fine-  
 326 tuning of continuous-time diffusion models as entropy-regularized control. *arXiv preprint  
 327 arXiv:2402.15194*, 2024.

328 [21] Siddarth Venkatraman, Mohsin Hasan, Minsu Kim, Luca Scimeca, Marcin Sendera, Yoshua  
 329 Bengio, Glen Berseth, and Nikolay Malkin. Outsourced diffusion sampling: Efficient posterior  
 330 inference in latent spaces of generative models. In *International Conference on Machine  
 331 Learning (ICML)*, 2025.

332 [22] Florentin Coeurdoux, Nicolas Dobigeon, and Pierre Chainais. Normalizing flow sampling with  
 333 langevin dynamics in the latent space. *arXiv preprint arXiv:2305.12149*, 2023.

334 [23] Matthias Schonlau, William J Welch, and Donald R Jones. Global versus local search in  
 335 constrained optimization of computer models. *Lecture notes-monograph series*, pages 11–25,  
 336 1998.

337 [24] Sebastian Ament, Samuel Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy.  
 338 Unexpected improvements to expected improvement for bayesian optimization. *Advances in  
 339 Neural Information Processing Systems*, 36:20577–20612, 2023.

340 [25] Nikolaus Hansen. The cma evolution strategy: a comparing review. *Towards a new evolutionary  
 341 computation: Advances in the estimation of distribution algorithms*, pages 75–102, 2006.

342 [26] Asma Atamna, Anne Auger, and Nikolaus Hansen. Augmented lagrangian constraint handling  
 343 for cma-es—case of a single linear constraint. In *International Conference on Parallel Problem  
 344 Solving from Nature*, pages 181–191. Springer, 2016.

345 [27] Siddarth Krishnamoorthy, Satvik Mehul Mashkaria, and Aditya Grover. Diffusion models for  
 346 black-box optimization. In *International Conference on Machine Learning (ICML)*, 2023.

347 [28] Dongxia Wu, Nikki Lijing Kuang, Ruijia Niu, Yian Ma, and Rose Yu. Diff-BBO: Diffusion-  
 348 based inverse modeling for black-box optimization. In *NeurIPS 2024 Workshop on Bayesian  
 349 Decision-making and Uncertainty*, 2024.

350 [29] Taeyoung Yun, Kiyoung Om, Jaewoo Lee, Sujin Yun, and Jinkyoo Park. Posterior inference with  
 351 diffusion models for high-dimensional black-box optimization. In *International Conference on  
 352 Machine Learning (ICML)*, 2025.

353 [30] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis.  
 354 *Advances in neural information processing systems*, 34:8780–8794, 2021.

355 [31] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medi-  
 356 cal imaging with score-based generative models. In *International Conference on Learning  
 357 Representations*, 2022.

358 [32] Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul  
 359 Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh Interna-  
 360 tional Conference on Learning Representations*, 2023.

361 [33] Ruiqi Feng, Chenglei Yu, Wenhao Deng, Peiyan Hu, and Tailin Wu. On the guidance of flow  
 362 matching. In *Forty-second International Conference on Machine Learning*, 2025.

363 [34] Cheng Lu, Huayu Chen, Jianfei Chen, Hang Su, Chongxuan Li, and Jun Zhu. Contrastive  
364 energy prediction for exact energy-guided diffusion sampling in offline reinforcement learning.  
365 In *International Conference on Machine Learning*, pages 22825–22855. PMLR, 2023.

366 [35] Luhuan Wu, Brian Trippe, Christian Naesseth, David Blei, and John P Cunningham. Practi-  
367 cal and asymptotically exact conditional sampling in diffusion models. *Advances in Neural*  
368 *Information Processing Systems*, 36:31372–31403, 2023.

369 [36] Gabriel Cardoso, Sylvain Le Corff, Eric Moulines, et al. Monte carlo guided denoising diffusion  
370 models for bayesian linear inverse problems. In *The Twelfth International Conference on*  
371 *Learning Representations*, 2024.

372 [37] Ying Fan, Olivia Watkins, Yuqing Du, Hao Liu, Moonkyung Ryu, Craig Boutilier, Pieter Abbeel,  
373 Mohammad Ghavamzadeh, Kangwook Lee, and Kimin Lee. Dpok: Reinforcement learning  
374 for fine-tuning text-to-image diffusion models. *Advances in Neural Information Processing*  
375 *Systems*, 36:79858–79885, 2023.

376 [38] Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion  
377 models with reinforcement learning. In *The Twelfth International Conference on Learning*  
378 *Representations*, 2024.

379 [39] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow  
380 matching for generative modeling. In *The Eleventh International Conference on Learning*  
381 *Representations*, 2023.

382 [40] Xingchao Liu, Chengyue Gong, et al. Flow straight and fast: Learning to generate and transfer  
383 data with rectified flow. In *The Eleventh International Conference on Learning Representations*,  
384 2023.

385 [41] Michael Samuel Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic  
386 interpolants. In *The Eleventh International Conference on Learning Representations*, 2023.

387 [42] Marcin Sendera, Minsu Kim, Sarthak Mittal, Pablo Lemos, Luca Scimeca, Jarrid Rector-Brooks,  
388 Alexandre Adam, Yoshua Bengio, and Nikolay Malkin. Improved off-policy training of diffusion  
389 samplers. *Advances in Neural Information Processing Systems*, 37:81016–81045, 2024.

390 [43] Nikolay Malkin, Moksh Jain, Emmanuel Bengio, Chen Sun, and Yoshua Bengio. Trajectory  
391 balance: Improved credit assignment in gflownets. *Advances in Neural Information Processing*  
392 *Systems*, 35:5955–5967, 2022.

393 [44] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and  
394 Ben Poole. Score-based generative modeling through stochastic differential equations. In  
395 *International Conference on Learning Representations*, 2021.

396 [45] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable  
397 predictive uncertainty estimation using deep ensembles. *Advances in neural information*  
398 *processing systems*, 30, 2017.

399 [46] Minsu Kim, Federico Berto, Sungsoo Ahn, and Jinkyoo Park. Bootstrapped training of score-  
400 conditioned generator for offline design of biological sequences. In *Advances in Neural*  
401 *Information Processing Systems (NeurIPS)*, 2023.

402 [47] Aviral Kumar and Sergey Levine. Model inversion networks for model-based optimization. In  
403 *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

404 [48] Ashvin Nair, Abhishek Gupta, Murtaza Dalal, and Sergey Levine. Awac: Accelerating online  
405 reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*, 2020.

406 [49] Zi Wang, Clement Gehring, Pushmeet Kohli, and Stefanie Jegelka. Batched large-scale bayesian  
407 optimization in high-dimensional spaces. In *International Conference on Artificial Intelligence*  
408 *and Statistics*, pages 745–754. PMLR, 2018.

409 [50] Kenan Šehić, Alexandre Gramfort, Joseph Salmon, and Luigi Nardi. Lassobench: A high-  
410 dimensional hyperparameter optimization benchmark suite for lasso. In *International Confer-  
411 ence on Automated Machine Learning*, pages 2–1. PMLR, 2022.

412 [51] José Miguel Hern, Michael A Gelbart, Ryan P Adams, Matthew W Hoffman, Zoubin Ghahra-  
 413 mani, et al. A general framework for constrained bayesian optimization using information-based  
 414 search. *Journal of Machine Learning Research*, 17(160):1–53, 2016.

415 [52] Nicola Demo, Marco Tezzele, and Gianluigi Rozza. A supervised learning approach involving  
 416 active subspaces for an efficient genetic algorithm in high-dimensional optimization problems.  
 417 *SIAM Journal on Scientific Computing*, 43(3):B831–B853, 2021.

418 [53] Linnan Wang, Rodrigo Fonseca, and Yuandong Tian. Learning search space partition for black-  
 419 box optimization using monte carlo tree search. *Advances in Neural Information Processing  
 420 Systems*, 33:19511–19522, 2020.

421 [54] Zeji Yi, Yunyue Wei, Chu Xin Cheng, Kaibo He, and Yanan Sui. Improving sample efficiency  
 422 of high dimensional bayesian optimization with mcmc. In *6th Annual Learning for Dynamics  
 423 & Control Conference*, pages 813–824. PMLR, 2024.

424 [55] Zelda B Zabinsky and Robert L Smith. Hit-and-run methods. *Encyclopedia of Operations  
 425 Research and Management Science*, pages 721–729, 2013.

426 [56] Leonard Papenmeier, Luigi Nardi, and Matthias Poloczek. Increasing the scope as you learn:  
 427 Adaptive bayesian optimization in nested subspaces. *Advances in Neural Information Processing  
 428 Systems*, 35:11586–11601, 2022.

429 [57] Nikolaus Hansen, Youhei Akimoto, and Petr Baudis. CMA-ES/pycma on Github. Zenodo,  
 430 DOI:10.5281/zenodo.2559634, February 2019.

431 [58] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint  
 432 arXiv:1606.08415*, 2016.

433 [59] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Inter-  
 434 national Conference on Learning Representations (ICLR)*, 2015.

435 [60] Yaron Lipman, Marton Havasi, Peter Holderrieth, Neta Shaul, Matt Le, Brian Karrer, Ricky TQ  
 436 Chen, David Lopez-Paz, Heli Ben-Hamu, and Itai Gat. Flow matching guide and code. *arXiv  
 437 preprint arXiv:2412.06264*, 2024.

438 [61] Ricky T. Q. Chen. torchdiffeq, 2018.

439 [62] Ulf Grenander and Michael I Miller. Representations of knowledge in complex systems. *Journal  
 440 of the Royal Statistical Society: Series B (Methodological)*, 56(4):549–581, 1994.

441 [63] Simon Duane, Anthony D Kennedy, Brian J Pendleton, and Duncan Roweth. Hybrid monte  
 442 carlo. *Physics letters B*, 195(2):216–222, 1987.

443 [64] J. H. Halton. Sequential monte carlo. *Mathematical Proceedings of the Cambridge Philosophical  
 444 Society*, 58(1):57–78, 1962.

445 [65] Nicolas Chopin. A sequential particle filter method for static models. *Biometrika*, 89(3):539–  
 446 552, 2002.

447 [66] John Skilling. Nested sampling for general bayesian computation. 2006.

448 [67] Pablo Lemos, Nikolay Malkin, Will Handley, Yoshua Bengio, Yashar Hezaveh, and Laurence  
 449 Perreault-Levasseur. Improving gradient-guided nested sampling for posterior inference. In  
 450 *International Conference on Machine Learning*, pages 27230–27253. PMLR, 2024.

451 [68] Qinsheng Zhang and Yongxin Chen. Path integral sampler: A stochastic control approach for  
 452 sampling. In *International Conference on Learning Representations*, 2022.

453 [69] Francisco Vargas, Will Sussman Grathwohl, and Arnaud Doucet. Denoising diffusion samplers.  
 454 In *The Eleventh International Conference on Learning Representations*, 2023.

455 [70] Julius Berner, Lorenz Richter, and Karen Ullrich. An optimal control perspective on diffusion-  
 456 based generative modeling. *Transactions on Machine Learning Research*, 2024.

457 [71] Lorenz Richter and Julius Berner. Improved sampling via learned diffusions. In *The Twelfth*  
458 *International Conference on Learning Representations*, 2024.

459 [72] Francisco Vargas, Shreyas Padhy, Denis Blessing, and Nikolas Nüsken. Transport meets varia-  
460 tional inference: Controlled monte carlo diffusions. In *The Twelfth International Conference on*  
461 *Learning Representations*, 2024.

462 [73] Salem Lahlou, Tristan Deleu, Pablo Lemos, Dinghuai Zhang, Alexandra Volokhova, Alex  
463 Hernández-García, Léna Néhale Ezzine, Yoshua Bengio, and Nikolay Malkin. A theory of  
464 continuous generative flow networks. In *International Conference on Machine Learning*, pages  
465 18269–18300. PMLR, 2023.

466 [74] Dinghuai Zhang, Ricky TQ Chen, Cheng-Hao Liu, Aaron Courville, and Yoshua Bengio.  
467 Diffusion generative flow samplers: Improving learning signals through partial trajectory  
468 optimization. In *The Twelfth International Conference on Learning Representations*, 2024.

469 [75] Nikolay Malkin, Salem Lahlou, Tristan Deleu, Xu Ji, Edward J Hu, Katie E Everett, Dinghuai  
470 Zhang, and Yoshua Bengio. GFlownets and variational inference. In *The Eleventh International*  
471 *Conference on Learning Representations*, 2023.

472 [76] Alexander Tong, Kilian FATRAS, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid  
473 Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based genera-  
474 tive models with minibatch optimal transport. *Transactions on Machine Learning Research*,  
475 2024.

476 [77] Aaron J Havens, Benjamin Kurt Miller, Bing Yan, Carles Domingo-Enrich, Anuroop Sriram,  
477 Daniel S. Levine, Brandon M Wood, Bin Hu, Brandon Amos, Brian Karrer, Xiang Fu, Guan-  
478 Horng Liu, and Ricky T. Q. Chen. Adjoint sampling: Highly scalable diffusion samplers via  
479 adjoint matching. In *Frontiers in Probabilistic Inference: Learning meets Sampling*, 2025.

480 [78] Minsu Kim, Sanghyeok Choi, Taeyoung Yun, Emmanuel Bengio, Leo Feng, Jarrid Rector-  
481 Brooks, Sungsoo Ahn, Jinkyoo Park, Nikolay Malkin, and Yoshua Bengio. Adaptive teachers for  
482 amortized samplers. In *The Thirteenth International Conference on Learning Representations*,  
483 2025.

484 **Appendix**

485 **A Task Details**

486 **A.1 Synthetic Functions**

We evaluate three synthetic functions in our constrained black-box optimization experiments: Rastrigin, Ackley, and Rosenbrock. The Rastrigin and Ackley functions are highly multi-modal functions with numerous local minima, whereas the Rosenbrock function features a narrow valley that makes convergence to the global minimum notoriously difficult [52]. Following [53, 54], we define the search domains as Rastrigin:  $[-5, 5]^D$ , Ackley:  $[-5, 10]^D$ , and Rosenbrock:  $[-5, 10]^D$ . All functions are subject to two constraints:

$$\sum_{d=1}^{200} x_d \leq 0 \quad \text{and} \quad \|\mathbf{x}\|_2^2 \leq 30$$

487 Although prior work enforced the tighter bound  $\|\mathbf{x}\|_2^2 \leq 5$ , we relax this constraint in our high-  
488 dimensional setting. For the indicator constraint experiments, we sample initial feasible points by  
489 hit-and-run MCMC [55].

490 **A.2 Rover Trajectory Planning**

Rover Trajectory Planning is a trajectory optimization task in a 2D environment introduced by [49]. The objective is to optimize the rover's trajectory, where its trajectory is represented by 30 points defining a B-Spline. We place 15 impassable obstacles  $o_i$  and impose collision-avoidance constraints  $c_i(\mathbf{x})$  as in [8]:

$$c_i(\mathbf{x}) = \begin{cases} -d(o_i, \gamma(\mathbf{x})) & \text{if } \gamma(\mathbf{x}) \cap o_i = \emptyset, \\ \max_{\alpha \in \gamma(\mathbf{x}) \cap o_i} \min_{\beta \in \partial o_i} d(\alpha, \beta) & \text{otherwise.} \end{cases}$$

491 where  $\gamma(\mathbf{x})$  denotes final trajectory,  $o_i$  is the region of the obstacle and  $\partial o_i$  denotes the boundary of  $o_i$ .  
492 A trajectory is feasible if and only if it does not intersect any obstacle. We follow the implementation  
493 from [49], but since there is no released code for the constraints, we implement the violation metric  
494 ourselves. Below is an example of the trajectory found by our method.




Figure 6: Trajectory found by CiBO, achieving regret of -4.59.

495 **A.3 Vehicle Design with 68 Constraints (MOPTA)**

496 MOPTA is the high-dimensional real-world problem of large-scale multidisciplinary mass optimiza-  
497 tion [6]. The objective is to minimize a vehicle’s mass, which incorporates decisions about materials,  
498 gauges, and vehicle shape with 68 performance constraints. The best-known optimum mass is  
499 approximately 222.74. We followed the implementation from [56].<sup>1</sup>

500 **A.4 LassoBench**

501 LassoBench [50]<sup>2</sup> is a high-dimensional benchmark for hyperparameter optimization, specifically  
502 designed to tune the hyperparameters of the Weighted LASSO (Least Absolute Shrinkage and  
503 Selection Operator) regression model. It includes both synthetic tasks (simple, medium, high, and  
504 hard) and real-world tasks (Breast cancer, Diabetes, Leukemia, DNA, and RCV1). In this work, we  
505 focus on the DNA task, a microbiology classification problem. It computes the average validation  
506 error across all cross-validation folds as an unconstrained objective. We reformulate the problem by  
507 retaining the mean validation error as the objective while introducing constraints that the validation  
508 error on each fold must not exceed 0.32.

509 **B Baselines Details**

510 In this section, we provide a thorough description of our baseline implementation details and specify  
511 the hyperparameter settings used across all experiments.

512 **DiBO** [29]: We use the original code<sup>3</sup>, and adapt DiBO to handle constrained optimization by  
513 reformulating the objective as a Lagrangian, setting the same  $\lambda$  value as our methods for fair  
514 comparison.

515 **DiffOPT** [15]: As there is no publicly available code, we re-implement this baseline on our own. To  
516 approximate the data distribution, we use diffusion models with a similar architecture to our method.  
517 To enable accurate sampling from the target distribution, we implement Langevin dynamics as the  
518 energy function, which can be constructed by surrogate models in our setting, is differentiable.

519 **DDOM** [27]: Building on the original implementation<sup>4</sup>, we reconstruct this baseline with network  
520 architecture matching our flow-based model. While maintaining the method’s specific parameters as  
521 specified in the original work, we incorporate a Lagrangian framework and set the same  $\lambda$  as ours.

522 **SCBO** [8]: We follow the tutorial code for SCBO provided by `botorch`<sup>5</sup> to reproduce the results.

523 **PCAGP-SCBO** [7]: To reproduce PCAGP-SCBO, we follow the code for SCBO and then apply  
524 `torch pca`<sup>6</sup> to project high-dimensional data into a reduced latent space with dimension  $l$  before  
525 fitting GP surrogates for constraints. For all synthetic tasks, we use  $l = 2$  and for real-world tasks,  
526 we conduct a hyperparameter search on  $[2, \lfloor D/2 \rfloor]$  and report the best one.

527 **cEI** [23]: We implement cEI acquisition function by using `qExpectedImprovement()` in `botorch`  
528 library. We train a GP surrogate model independently for the objective and each constraint.

529 **LogcEI** [24]: We implement logcEI acquisition function by using `qLogExpectedImprovement()` in  
530 `botorch` library. We train a GP surrogate model independently for the objective and each constraint.

531 **CMA-ES** [25]: We employ the `pycma`<sup>7</sup> library [57]. For constraint handling, we formulate the  
532 problem using the same Lagrangian approach with the same  $\lambda$  value as ours for each task.

---

<sup>1</sup><https://github.com/LeoIV/BAxUS>

<sup>2</sup><https://github.com/ksehic/LassoBench>

<sup>3</sup><https://github.com/umkiyoung/DiBO>

<sup>4</sup><https://github.com/siddarthk97/ddom>

<sup>5</sup>[https://botorch.org/docs/tutorials/scalable\\_constrained\\_bo/](https://botorch.org/docs/tutorials/scalable_constrained_bo/)

<sup>6</sup>[https://github.com/valentingol/torch\\_pca](https://github.com/valentingol/torch_pca)

<sup>7</sup><https://github.com/CMA-ES/pycma>

**Algorithm 1** CiBO

---

1: **Input:** Initial dataset  $\mathcal{D}_0$ ; Max rounds  $R$ ; Batch size  $B$ ; Buffer size  $L$ ; Number of constraints  $M$ ;  
     Flow model  $p_\theta$ ; Diffusion sampler  $p_\psi$ ; Proxies  $f_{\phi_1}, \dots, f_{\phi_K}, g_\phi^{(1)}, \dots, g_\phi^{(M)}$ ;  
 2: **for**  $r = 0, \dots, R - 1$  **do**  
 3:     Initialize  $p_\theta, p_\psi, f_{\phi_1}, \dots, f_{\phi_K}, g_\phi^{(1)}, \dots, g_\phi^{(M)}$   
 4:  
 5:     **Phase 1. Training Models**  
 6:     Compute weights  $w(y, \mathbf{c}, \mathcal{D}_r)$  with Equation (7)  
 7:     Train  $p_\theta$  with Equation (8)  
 8:     Train  $f_{\phi_1}, \dots, f_{\phi_K}, g_\phi^{(1)}, \dots, g_\phi^{(M)}$  with Equation (9)  
 9:  
 10:     **Phase 2. Sampling Candidates**  
 11:     Train  $p_\psi$  with Equation (14) using prior  $p_\theta$  and  $\mathbf{z} \sim N(0, \mathbf{I})$   
 12:     Sample latent noise with  $\{\mathbf{z}_i\}_{i=1}^{NB} \sim p_\psi(\mathbf{z})$   
 13:     Projection to data space with learned mapping  $\mathbf{x}_i = f_\theta(\mathbf{z}_i) \quad \forall i \in \{1, \dots, NB\}$   
 14:  
 15:     **Filtering**  
 16:     Select top- $B$  samples  $\{\mathbf{x}_b\}_{b=1}^B$  with respect to:  
          $r_\phi(\mathbf{x}_i) - \lambda \sum_{m=1}^M \max(0, g_\phi^{(m)}(\mathbf{x}_i)) \quad \forall i = \{1, \dots, NB\}$   
 17:  
 18:     **Evaluation and Moving Dataset**  
 19:     Evaluate  $y_b = f(\mathbf{x}_b), \quad c_b^m = g^{(m)}(\mathbf{x}_b) \quad \forall m = \{1, \dots, M\} \quad \forall b = \{1, \dots, B\}$   
 20:     Update  $\mathcal{D}_{r+1} \leftarrow \mathcal{D}_r \cup \{(\mathbf{x}_b, y_b, \mathbf{c}_b)\}_{b=1}^B$   
 21:     **if**  $|\mathcal{D}_{r+1}| > L$  **then**  
 22:         Remove last  $|\mathcal{D}_{r+1}| - L$  samples from  $\mathcal{D}_{r+1}$  with respect to:  $y - \lambda \sum_{m=1}^M \max(0, c^m)$   
 23:     **end if**  
 24: **end for**

---

534 **D Implementation Details**

535 In this section, we introduce the implementation details of our method **CiBO**. Specifically, model  
536 architectures, the training processes employed, the hyperparameter configurations used, and the  
537 computational resources required.

538 **D.1 Training Models**

539 **D.1.1 Training Proxies**

540 We employ an ensemble of five proxies to model the objective function and a single proxy for each  
541 constraint. Each proxy is implemented as a MLP with three hidden layers of 1024 units, using GELU  
542 [58] activations. Proxies are trained with the Adam optimizer [59] for 100 epochs per round at a  
543 learning rate of  $1 \times 10^{-3}$  and a batch size of 256. All hyperparameters related to the proxy are listed  
544 in Table 1.

Table 1: Hyperparameters for Training Proxy

	Parameters	Values
Architecture	Num Ensembles	5
	Number of Layers	3
	Num Units	1024
Training	Batch size	256
	Optimizer	Adam
	Learning Rate	$1 \times 10^{-3}$
	Training Epochs	100

545 **D.1.2 Training Flow-based Models**

546 We adopt the architecture of [60] for our flow model, comprising three hidden layers with 512  
547 units each. Training is performed using Adam optimizer for 500 epochs per round, with a learning  
548 rate of  $1 \times 10^{-3}$  and a batch size of 256. For ODE integration during sampling, we employ the  
549 Runge-Kutta 4 method with `torchdiffeq` [61], and set the integration steps as 250. All flow-model  
hyperparameters are detailed in Table 2.

Table 2: Hyperparameters for Training Flow-based Model

	Parameters	Values
Architecture	Number of Layers	3
	Num Units	512
Training	Batch size	256
	Optimizer	Adam
	Learning Rate	$1 \times 10^{-3}$
	Training Epochs	500

551 **D.2 Sampling Candidates**

552 **D.2.1 Training Diffusion Sampler**

553 Various approaches have been developed to draw samples from a distribution when only an unnormalized probability density or energy function is available. Traditional methods include Markov  
554 Chain Monte Carlo (MCMC) techniques [62, 63, 64, 65, 66, 67], though their computational cost  
555 increases dramatically in high-dimensional spaces. More recently, amortized variational inference  
556 methods, particularly those based on training diffusion samplers [68, 69, 70, 71, 72, 73, 74], have  
557 gained widespread adoption as they offer improved scalability for high-dimensional problems.

559 Following the [21], we adopt [42] to train diffusion sampler to sample from the target:

$$p_{\text{post}}(\mathbf{z}) \propto p(\mathbf{z}) \exp \left( \beta \cdot \left[ r_{\phi}(f_{\theta}(\mathbf{z})) - \lambda \sum_{m=1}^M \max \left( 0, g_{\phi}^{(m)}(f_{\theta}(\mathbf{z})) \right) \right] \right) \quad (15)$$

560 Here, the right-hand-side term serves as an unnormalized probability density, which the diffusion  
561 sampler amortizes the sampling cost by approximating it.

**Off-policy Training of Diffusion Sampler** As mentioned in the Section 4.2, we use the Trajectory Balance objective to train the diffusion sampler.

$$\mathcal{L}_{\text{TB}}(\mathbf{z}_{0:1}; \psi) = \left( \log \frac{Z_{\psi} p(\mathbf{z}_0) \prod_{i=0}^{T-1} p_F(\mathbf{z}_{(i+1)\Delta t} | \mathbf{z}_{i\Delta t}; \psi)}{p(\mathbf{z}_1) r(f_{\theta}(\mathbf{z}_1)) \prod_{i=1}^T p_B(\mathbf{z}_{(i-1)\Delta t} | \mathbf{z}_{i\Delta t})} \right)^2$$

562 The primary advantage of the TB loss is off-policy training [42, 75]. We can train our model not only  
563 from the on-policy trajectories through the reverse SDE  $\{\mathbf{z}_0, \dots, \mathbf{z}_1\} = \tau \sim p_F(\tau)$  but also from  
564 the trajectories through the forward SDE conditioned on the generated samples  $\tau \sim p_B(\tau | \mathbf{z}_1)$ . This  
565 proves its effectiveness on mode coverage and credit assignment [42].

566 Specifically, we repeat two processes. First, we sample trajectories on-policy  $\tau \sim p_F(\tau)$ , train  
567 the model with Equation (14), and collect the samples  $\mathbf{z}_1$  into the buffer. Second, from the col-  
568 lected samples  $\mathbf{z}_1$ , we generate off-policy trajectories through  $\tau \sim p_B(\tau | \mathbf{z}_1)$ , then train with  
569 the Equation (14). During the off-policy training, we prioritize the samples with low energy:  
570  $\mathcal{E}(\mathbf{z}_1) = -\log(p(\mathbf{z}_1)r(f_{\theta}(\mathbf{z}_1)))$  following [42] to make our model focus on the low energy samples.  
571 These techniques improve the overall performance of our framework (Appendix E.4).

572 We use the original code<sup>8</sup> released from [42] for implementation. We also set method-specific  
573 hyperparameters with Path Integral Sampler (PIS) [68] architecture, zero initialization, and t-scale  
574 to 1 to make sure the initialized  $p_F(\mathbf{z}_1)$  starts from the standard normal distribution. Detailed  
575 hyperparameters for training the diffusion sampler can be found in Table 3.

Table 3: Hyperparameters for Training Diffusion Sampler

	Parameters	Values
Architecture	Number of Layers	2
	Num Units	256
	Diffusion Time Steps	50
Training	Batch size	256
	Optimizer	Adam
	Learning Rate	$1 \times 10^{-3}$
	Training Epochs	50

576 **Computational Resources.** Our experiments were conducted using NVIDIA RTX 3090 and A6000  
577 GPUs. These resources were sufficient to train our models within a reasonable time for all reported  
578 experiments. Details of computational time can be found at Appendix F.

<sup>8</sup><https://github.com/GFNOrg/gfn-diffusion>

579 **D.3 Hyperparameters**

580 In our formulation of constrained black-box problems, we introduce  $\lambda$  for Lagrangian augmentation.  
 581 We draw  $N \times B$  samples from the posterior distribution, then select  $B$  samples during filtering. After  
 582 evaluation, we update the training set by keeping the top  $L$  highest-scoring samples subject to the  
 583 Lagrangian objective. Table 4 summarizes all hyperparameter values used in candidate selection,  
 584 and we include additional analysis to assess how each parameter affects overall performance in  
 585 Section 5.4 and Appendix E.

Table 4: Hyperparameters during sampling candidates

	Lambda $\lambda$	Inverse Temperature $\beta$	Buffer Size $L$	Filtering Coefficient $N$
Ackley 200D	10	$10^5$	3000	10
Rastrigin 200D	10	$10^5$	2000	10
Rosenbrock 200D	10	$10^5$	2000	10
RoverPlanning 60D	3	$10^5$	1000	10
Mopta 124D	3	$10^3$	500	10
DNA 180D	5	$10^3$	1000	15

586 **E Further Analysis**

587 In this section, we provide further analysis on different components of our method that are not  
 588 included in the main manuscript due to the page limit.

589 **E.1 Analysis on Feasibility Ratio**

590 To further analyze our method’s ability to effectively handle constraints, we report the feasibility  
 591 ratio across optimization batches for the Rastrigin 200D task. Here, the feasibility ratio denotes the  
 592 number of feasible samples over queried samples.




Figure 7: Feasibility ratio over all baselines. Experiments are conducted with four random seeds, and the mean and one standard deviation are reported.

593 As shown in Figure 7, CiBO demonstrates superior performance by rapidly achieving the highest  
 594 feasibility ratio within the first 5-10 batches, significantly faster than all competing methods. While  
 595 some baselines (SCBO, PCAGP-SCBO) eventually reach high feasibility ratios, they require approxi-  
 596 mately twice as many batches to achieve comparable performance. Other methods like DiBO and  
 597 DiffOPT take even longer (around 20 batches), and CMA-ES struggles substantially, only reaching  
 598 moderate feasibility ratios after 50 batches. Notably, CiBO not only reaches the high feasibility ratio  
 599 faster but also maintains it consistently throughout the optimization process, demonstrating its robust  
 600 constraint-handling capabilities in high-dimensional spaces.

601 **E.2 Analysis on Filtering coefficient  $N$**

602 To improve the sample efficiency of our method, we introduce filtering, where we sample  $N \times B$   
 603 candidates from the posterior distribution, then select the highest  $B$  samples with respect to the  
 604 Lagrangian-relaxed objective function. To analyze the impact of the filtering coefficient  $N$ , we  
 605 experiment with varying  $N$  values, including our default  $N = 100$ .




Figure 8: Performance of CiBO in Rastrigin-200D and Rover Planning-60D with varying  $N$ . Experiments are conducted with four random seeds, and the mean and one standard deviation are reported.

606 As shown in Figure 8, increasing the filtering coefficient improves sample efficiency by concentrating  
 607 candidate selection in both high objective values and feasible regions. If the coefficient is set too low,  
 608 we lose its exploitation capability, leading to slower convergence.

609 **E.3 Analysis on Buffer Size  $L$**

610 In each round, we retain the  $L$  top-scoring samples with respect to the Lagrangian-relaxed objective  
 611 function for computational efficiency. To analyze the effect of the buffer size  $L$ , we conduct  
 612 experiments by varying  $L$ . As demonstrated in Figure 9, using too small  $L$  occasionally gets stuck in  
 613 a sub-optimal solution while using too large  $L$  exhibits a slow convergence rate.




Figure 9: Performance of CiBO in Rastrigin-200D and Rover Planning-60D with varying  $L$ . Experiments are conducted with four random seeds, and the mean and one standard deviation are reported.

614 **E.4 Effect of Off-policy Training in Amortized Inference**

615 We employ off-policy training with the TB loss to train the diffusion sampler as detailed in Section 4.2.  
 616 To analyze the impact of off-policy training on performance, we conduct ablation studies on different  
 617 training schemes. As shown in Figure 10, off-policy training consistently outperforms on-policy  
 618 methods, and the performance gap widens as the number of constraints grows (15 constraints in  
 619 Rover Planning versus only 2 in Rastrigin). It underlines that training with off-policy samples is  
 620 crucial for amortizing the posterior distribution with multiple modes and a large plateau.




Figure 10: Comparison between off-policy and on-policy in Rastrigin-200D and Rover Planning-60D. Experiments are conducted with four random seeds, and the mean and one standard deviation are reported.

621 **E.5 Analysis on Initial Dataset size  $|D_0|$  and Batch size  $B$**

622 The size of the initial dataset,  $|D_0|$ , and batch size  $B$  play a critical role in the performance of  
 623 black-box optimization algorithms. When  $|D_0|$  is small and  $B$  is large, the algorithm must optimize  
 624 using very limited information, making the search significantly more challenging. To this end, we  
 625 conduct experiments varying  $|D_0|$  and  $B$  to demonstrate the robustness of our method on initial data  
 626 configurations. As shown in Figure 11, our method demonstrates robustness regarding both the initial  
 627 dataset size  $|D_0|$  and the batch size  $B$ .




Figure 11: Performance of CiBO in Rastrigin-200D with varying  $|D_0|$  and  $B$ . Experiments are conducted with four random seeds, and the mean and one standard deviation are reported.

628 **F Runtimes**

629 We report the running time of each method in Table 5. To measure the runtime, we conduct  
 630 experiments on a single NVIDIA RTX 3090 GPU and Intel Xeon Platinum CPU @ 2.90 GHz.  
 631 As shown in the table, the running time of our method is similar to other generative model-based  
 approaches, and mostly faster than BO-based methods.

Table 5: Average time (in seconds) for each round in each method.

	Rastrigin-200D	Ackley-200D	Rosenbrock-200D	RoverPlanning-60D	Mopta-124D	DNA-180D
cEI	336.96 $\pm$ 47.48	133.12 $\pm$ 6.66	489.86 $\pm$ 81.38	111.13 $\pm$ 5.83	205.66 $\pm$ 5.35	133.98 $\pm$ 9.16
LogcEI	720.45 $\pm$ 56.76	158.38 $\pm$ 10.13	593.39 $\pm$ 97.93	113.28 $\pm$ 3.97	324.27 $\pm$ 9.91	161.08 $\pm$ 8.21
SCBO	322.81 $\pm$ 47.38	117.81 $\pm$ 6.50	475.02 $\pm$ 80.11	87.83 $\pm$ 4.42	270.30 $\pm$ 5.19	147.83 $\pm$ 12.57
PCAGP-SCBO	327.48 $\pm$ 51.20	122.67 $\pm$ 3.50	479.06 $\pm$ 82.41	17.55 $\pm$ 3.05	20.69 $\pm$ 0.39	81.34 $\pm$ 9.19
CMA-ES	0.08 $\pm$ 0.00	0.09 $\pm$ 0.00	0.10 $\pm$ 0.01	0.61 $\pm$ 0.00	5.33 $\pm$ 0.15	46.58 $\pm$ 3.16
DDOM	26.87 $\pm$ 0.28	27.00 $\pm$ 0.32	26.96 $\pm$ 0.12	3.56 $\pm$ 0.02	8.63 $\pm$ 0.37	50.99 $\pm$ 0.81
DifFOPT	91.00 $\pm$ 5.07	111.04 $\pm$ 1.51	89.37 $\pm$ 8.27	29.48 $\pm$ 1.21	105.61 $\pm$ 2.71	60.64 $\pm$ 1.34
DiBO	73.97 $\pm$ 0.56	68.89 $\pm$ 0.98	73.83 $\pm$ 1.11	29.43 $\pm$ 1.51	66.61 $\pm$ 4.16	71.85 $\pm$ 2.48
<b>CiBO</b>	73.39 $\pm$ 2.09	103.43 $\pm$ 4.84	82.24 $\pm$ 6.50	53.58 $\pm$ 5.14	105.43 $\pm$ 2.24	81.77 $\pm$ 2.37

632

633 **G Limitations and Future Work**

634 We are interested in improving our method further. First, as we need to train all models with the  
635 updated dataset in every round, presenting a framework that can efficiently reuse the trained models  
636 from the previous rounds would be beneficial. Furthermore, there are several advancements in the  
637 literature on flow-based model training [76] and diffusion samplers [77, 78], which could potentially  
638 yield substantial performance gains. We leave them as future work.

639 **H Broader Impact**

640 Advances in real-world design optimization have the potential to drive major innovations, but they  
641 also come with potential risks and unintended consequences. For example, optimization techniques  
642 in biochemical design may uncover novel compounds with therapeutic potential, but similar methods  
643 could also be misused to discover harmful substances. It is essential for researchers to act responsibly  
644 and ensure their work serves the public good.