

Nesterov smoothing for sampling without smoothness

Jiaojiao Fan, Bo Yuan, Jiaming Liang and Yongxin Chen

Abstract—We study the problem of sampling from a target distribution in \mathbb{R}^d whose potential is not smooth. Compared with the sampling problem with smooth potentials, this problem is much less well-understood due to the lack of smoothness. In this paper, we propose a novel sampling algorithm for a class of non-smooth potentials by first approximating them by smooth potentials using a technique that is akin to Nesterov smoothing. We then utilize sampling algorithms on the smooth potentials to generate approximate samples from the original non-smooth potentials. We select an appropriate smoothing intensity to ensure that the distance between the smoothed and un-smoothed distributions is minimal, thereby guaranteeing the algorithm’s accuracy. Hence we obtain non-asymptotic convergence results based on existing analysis of smooth sampling. We verify our convergence result on a synthetic example and apply our method to improve the worst-case performance of Bayesian inference on a real-world example.

I. INTRODUCTION

Sampling from a target distribution $\pi(x) \propto \exp(-s(x))$ known up to a normalization constant is an important problem in many areas such as (particle) filtering and estimation, inverse problems, machine learning, and plays a pivotal role in Bayesian statistics and inference [1], [2], [3], [4]. For instance, in nonlinear filtering problem, one can infer from the posterior distribution of the state by sampling from it. Over the past decades, there has been a vast amount of research carried out on sampling problems from smooth distributions [5], [6], i.e., the gradient of the potential $s(x)$ is Lipschitz. In contrast, the understanding of non-smooth sampling is still relatively limited.

In this work, we consider the task of sampling from a specific class of non-smooth distributions

$$\pi(x) \propto \exp(-s(x))$$

whose potentials $s(x)$ permit an explicit max-structure as

$$s(x) = f(x) + \max_{y \in \mathcal{Y}} \{h(x, y) - g(y)\} \quad (1)$$

where f is smooth, g is convex, \mathcal{Y} is convex and bounded, and h is Lipschitz and smooth. $s(x)$ can be non-smooth. For instance, when $f(x) = 0$ and $h(x) = Ax$, $s(x)$ becomes a piece-wise affine function and is clearly non-smooth. A special instance of particular interest is

$$s(x) = f(x) + \max_i \{h_1(x), h_2(x), \dots, h_n(x)\}. \quad (2)$$

This work was supported by the NSF under grant 1942523 and 2008513.

J. Fan, B. Yuan and Y. Chen are with the School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA {jiaojiaofan, byuan48, yongchen}@gatech.edu

J. Liang is with the Department of Computer Science, Yale University, New Haven, CT 06520, USA jiaming.liang@yale.edu

The optimization of minimizing $s(x)$ has been widely used for science and engineering problems to optimize the worst-case performance [7], [8]. Its Bayesian counterpart that accounts for uncertainties can be captured by a sampling problem from the potential $s(x)$.

Compared to other recent work [9], [10] for sampling from black box non-smooth distribution, we assume certain structures in the target distribution and hope to leverage these structures to achieve better performance. Indeed, in practice, it is rare to have a black box model as a target distribution; we almost always know some structures of the problem. It could be advantageous to utilize such structures in the algorithm design. With this in mind, we take one step forward in the direction of sampling beyond black box models and propose a sampling method for non-smooth distribution that utilizes the specific max-structure (1).

Our strategy is akin to Nesterov smoothing [11] in optimization that constructs a smooth function $s_\beta(x)$ whose difference with $s(x)$ can be controlled through the smoothing intensity β . In Nesterov smoothing, the idea is to use a fast smooth optimization algorithm to minimize $s_\beta(x)$ for the purpose of optimizing the non-smooth function $s(x)$. In our problem, instead, we sample from $\pi_\beta(x) \propto \exp(-s_\beta(x))$ for the purpose of sampling from $\pi(x) \propto \exp(-s(x))$. Thanks to the smoothness of $s_\beta(x)$, we can take advantage of the many existing sampling algorithms for smooth potentials to sample from $\pi_\beta(x) \propto \exp(-s_\beta(x))$. Our smoothing strategy is compatible with any such algorithms requiring only first-order smoothness. The complexity bounds we obtained with Nesterov smoothing have exactly the same dimension dependency as their smooth counterparts. In a high level, this work is along the direction of the recent line of research that tries to bridge optimization and sampling [12], [13], [14]. Our results prove that some smoothing techniques used in optimization are equally effective in sampling.

a) Related Works: Over the last few years, several algorithms for sampling from potentials that are not smooth have been developed. A majority of them is devoted to the analysis of Unadjusted Langevin Monte Carlo (LMC) in the non-smooth setting [15], [16], [10]. These theoretical results are challenging and significantly different from those in the smooth setting. In [17], [18], [19], the non-smooth sampling problem is studied and analyzed from an optimization point of view. Another algorithm that is effective for non-smooth sampling is the proximal sampler developed in the sequence of work [20], [9], [14], [21]. The methods that are most related to ours are Gaussian smoothing [22], and Moreau envelope [23], [24]. Gaussian smoothing can be applied to any convex potential, however the evaluation of the smoothed potential is

difficult. Normally one can only get unbiased estimation of it through Gaussian sampling, inducing additional variance in the algorithm. Moreau envelope can be applied to any weakly-convex potential. To evaluate the smoothed potential or its gradient, one needs to solve an optimization problem (to get the proximal map), inducing additional complexity, especially when the potential is not convex. We include a more detailed comparison with these two smoothing techniques in the end of Section III-B.

We summarize our contributions below:

- i) Inspired by Nesterov smoothing, we develop a smoothing technique, different from the existing ones like Gaussian smoothing and Moreau envelope, for sampling problems for a class of potentials that permit the max-structure (1). Our smoothing technique is compatible with any first-order sampling algorithm for smooth potentials and the convergence is guaranteed via a proper smoothing intensity β .
- ii) Combining this smoothing technique with existing sampling algorithms for smooth potentials we obtain non-asymptotic complexity bounds for strongly-log-concave distributions, log-concave distributions, and distributions satisfying log-Sobolev inequality (LSI). In particular, the complexity has the same dimensional dependency as its smooth counterparts. Besides, our analysis covers certain composite potentials with non-smooth (Lipschitz continuous) components and the distribution satisfies LSI.
- iii) We show one synthetic example to verify our non-asymptotic analysis and demonstrate that the structure (1) can be useful for improving the worst-case performance through a real-world logistic regression example.

II. BACKGROUND

We assume the norm in the space of x and y can be different and denote them as $\|\cdot\|_{\mathcal{X}}$ and $\|\cdot\|_{\mathcal{Y}}$ respectively. We use $J_h(x) : \mathbb{R}^d \rightarrow \mathbb{R}^{n \times d}$ to represent the Jacobian matrix of the function h at x . The norm of a matrix $A : \mathbb{R}^d \rightarrow \mathbb{R}^{n \times d}$ is defined as $\|A\|_{\mathcal{X}, \mathcal{Y}} = \max_{x, y} \{ \langle Ax, y \rangle_{\mathcal{Y}} : \|x\|_{\mathcal{X}} = 1, \|y\|_{\mathcal{Y}} = 1 \}$. Given a primal space equipped with norm $\|\cdot\|$, we denote the norm in its dual space as $\|\cdot\|^*$. In this article, we assume the probability measure have density with respect to the Lebesgue measure.

A. Nesterov smoothing

It is well-known that with subgradient methods, the optimal bound of non-smooth convex minimization complexity is $\mathcal{O}(1/\varepsilon^2)$ to achieve ε error in the function value [25], [26]. However, beyond the black-box oracle model, there are various structures that can be used to improve this bound. Nesterov smoothing [11] is an interesting technique for functions with explicit max structure (1). In doing so, the complexity order can be improved to $\mathcal{O}(1/\varepsilon)$. The underlying principle of Nesterov smoothing is that *the dual of a strongly-convex function is smooth and vice versa* [27]. Specifically, it adds a σ -strongly-convex penalty function $\ell(y)$ to the original maximization such that

$$s_{\beta}(x) = f(x) + \max_{y \in \mathcal{Y}} \{ \langle h(x), y \rangle - g(y) - \beta \ell(y) \}, \quad (3)$$

where β is the smoothing intensity. In the original paper of Nesterov smoothing, they consider the specific task $h(x) = Ax$ and find s_{β} is smooth with constant $L_f + \|A\|_{\mathcal{X}, \mathcal{Y}}^2 / \beta \sigma$. Then they apply the Nesterov accelerated gradient descent algorithm to minimize s_{β} , an optimal scheme for smooth optimization. Our idea is very similar on the high level: use smoothing for the non-smooth potential first, and then apply the existing sampling algorithm for the smoothed distribution.

B. Smooth sampling

In this section, we summarize several popular sampling algorithms for smooth distributions. In general, there are several dominant popular algorithms: LMC [28], Kinetic Langevin Monte Carlo (KLMC) [29] Metropolis-Adjusted Langevin Algorithm (MALA) [30], and Hamiltonian Monte Carlo (HMC) [31]. Suppose we are given the target distribution $\pi \propto \exp(-V(x))$ for some smooth potential $V(x)$. The LMC algorithm runs the update

$$x_{k+1} = x_k - \gamma \nabla V(x_k) + \sqrt{2\gamma} \varepsilon_k, \quad (4)$$

where ε_k is the standard Gaussian random variable in \mathbb{R}^d . The initial point x_0 is user-defined and the step size $\gamma > 0$ is chosen sufficiently small to ensure limited discretization error. Indeed, (4) is the Euler discretization of Langevin diffusion dynamics,

$$dx_t = -\nabla V(x_t) dt + \sqrt{2} dB_t. \quad (5)$$

This dynamics has an invariant density π no matter what x_0 is. The underdamped Langevin diffusion adds a Hamiltonian ingredient to the standard Langevin diffusion (5), so that an additional random variable $v_t := dx_t/dt$ controls the velocity in the dynamics. Similarly, HMC maintains a velocity term, but with a different randomization manner. MALA is also a variant of LMC. As its name implies, it adds a Metropolis-Hastings rejection step to LMC, in order to eliminate the bias induced by discretization.

In spite of the canonical role of LMC in sampling, it is not the fastest algorithm empirically. Theoretically, in the standard strongly-convex potential case, the complexity of LMC can achieve $\tilde{\mathcal{O}}(d)$ [5], whereas KLMC, HMC algorithms can achieve $\tilde{\mathcal{O}}(d^{1/2})$ [32], [33]. The complexity $\tilde{\mathcal{O}}(d^{1/2})$ also holds for MALA, however, under a warm start [34]. Decades of extensive research have led to the recent proposal of several innovative sampling algorithms and their variants. Notable examples include the proximal algorithm [20], [35] based on the restricted Gaussian oracle and the randomized midpoint method for KLMC [36]. The proximal algorithm has been proved to have the state-of-art dimension dependency for log-concave distribution [35].

III. NESTEROV SMOOTHING FOR SAMPLING

In this section, we will introduce our Nesterov smoothing sampling algorithm in Section III-A and bound the error caused by smoothing in Section III-B.

A. Problem setup and algorithm

Formally, we consider the task to sample from a non-smooth distribution $\pi(x) \propto \exp(-s(x))$ where

$$s(x) = f(x) + \max_{y \in \mathcal{Y}} \{ \langle h(x), y \rangle - g(y) \}.$$

We will assume $f(x) : \mathcal{X} \rightarrow \mathbb{R}$ is L_f -smooth; $h(x) : \mathcal{X} \rightarrow \mathcal{Y}$ is λ_h -Lipschitz and L_h -smooth; $g(y) : \mathcal{Y} \rightarrow \mathbb{R}$ is continuous and convex through out the paper. A function $f(x)$ is L -smooth if

$$\|\nabla f(x_1) - \nabla f(x_2)\|_{\mathcal{X}}^* \leq L\|x_1 - x_2\|_{\mathcal{X}}$$

for $\forall x_1, x_2 \in \mathbb{R}^d$. Correspondingly, we say a multivariate function $h(x)$ is L -smooth if

$$\|J_h(x_1) - J_h(x_2)\|_{\mathcal{X}, \mathcal{Y}} \leq L\|x_1 - x_2\|_{\mathcal{X}}$$

for $\forall x_1, x_2 \in \mathcal{X}$. We say a multivariate function $h(x)$ is λ -Lipschitz if

$$\|h(x_1) - h(x_2)\|_{\mathcal{Y}}^* \leq \lambda\|x_1 - x_2\|_{\mathcal{X}},$$

which is equivalent to $\|J_h(x)\|_{\mathcal{X}, \mathcal{Y}} \leq \lambda$ for $\forall x \in \mathcal{X}$ if $h(x)$ is differentiable. This boils down to that the operator norm of A is bounded when $h(x) = Ax$; this allows us to compare the smoothing effect with the case considered in [11]. Without more specifications, we assume $\|\cdot\|_{\mathcal{X}}^* = \|\cdot\|_{\mathcal{X}} = \|\cdot\|_2$ is Euclidean norm to be consistent with the sampling literature.

It is very likely that $s(x)$ is not differentiable since it is a "piece-wise" function segmented by different maximizer y . As such, it is not feasible to directly apply the smooth sampling algorithms. Inspired by the Nesterov smoothing technique, our strategy is to 1) smooth $s(x)$ as $s_\beta(x)$, 2) apply smooth sampling algorithm to $\pi_\beta(x)$.

For smoothing, we consider the surrogate function in equation (3) where $\ell(y)$ is a σ -strongly-convex prox-function

$$\ell(y) \geq \frac{\sigma}{2} \|y - y_0\|_{\mathcal{Y}}^2.$$

The center of \mathcal{Y} is defined as $y_0 = \arg \min_y \{ \ell(y) : y \in \mathcal{Y} \}$ and we can assume that $\ell(y_0) = 0$. The compactness of \mathcal{Y} also allows us to define

$$D = \max_y \{ \ell(y) : y \in \mathcal{Y} \}, \quad R = \max_y \{ \|y\|_{\mathcal{Y}} : y \in \mathcal{Y} \}.$$

From the definition, D is the squared diameter of the set \mathcal{Y} , and R is the furthest distance of any point in the set to the origin. By the triangular inequality, it holds that

$$R \leq \|y_0\| + \sqrt{\frac{2D}{\sigma}}.$$

In this way, $s_\beta(x)$ is a smooth function, which is justified by Lemma 1. Proofs can be found in the extended version [37].

Lemma 1: Let $\phi_\beta(x, y) := \langle h(x), y \rangle - g(y) - \beta \ell(y)$, and $y_\beta(x) := \arg \max_{y \in \mathcal{Y}} \phi_\beta(x, y)$, then

$$\nabla s_\beta(x) = \nabla f(x) + J_h^\top(x) y_\beta(x). \quad (6)$$

Moreover, $s_\beta(x)$ is L_{s_β} -smooth with

$$L_{s_\beta} = L_f + RL_h + \frac{\lambda_h^2}{\beta \sigma}.$$

Algorithm 1 Nesterov smoothing sampling framework

Input: A first-order based smooth sampling algorithm, the target distribution $\pi \propto \exp(-s(x))$, number of iterations K , initial point $x_0 \sim \mu_0$

Initialization: Choose the smoothing intensity β

for $k = 0, \dots, K - 1$ **do**

Calculate $\nabla s_\beta(x)$ according to (6)

Iterate x_{k+1} by the smooth sampling algorithm

end for

Output: x_K

If $h(x) = Ax$, then $L_h = 0$ and $\lambda_h = \|A\|_{\mathcal{X}, \mathcal{Y}}$, we get $L_{s_\beta} = L_f + \frac{1}{\beta \sigma} \|A\|_{\mathcal{X}, \mathcal{Y}}^2$ in this special case. This recovers the smoothness constant in [11, Theorem 1]. In many cases, λ_h is dimensional-free. For example, when $h(x) = Ax$ and $A = cI_m$, then $\lambda_h = c$ regardless of the dimensionality.

We will also assume that we can solve $\max_y \phi_\beta(x, y)$ for any $x \in \mathcal{X}$ in a fast dimension-free manner, which has been verified by the complexity analysis [38, Section 3] and especially several examples in [11, Section 4]. Thus we can query the gradient of $s_\beta(x)$ according to (6) and resort to first-order sampling algorithm to sample from $\pi_\beta \propto \exp(-s_\beta(x))$. The pseudo-code of our scheme is presented in Algorithm 1.

B. Bounded distance between π and π_β

When $\beta \rightarrow 0$, it holds that $\pi_\beta \rightarrow \pi$. In this section, we rigorously bound the distance between π and π_β . A key observation is that $s_\beta(x)$ is a uniform smooth approximation of $s(x)$ [11, Eq. 2.7], i.e.,

$$s_\beta(x) \leq s(x) \leq s_\beta(x) + \beta D.$$

To bound the difference between π and π_β , we select two widely-used distributional distances as the error criteria. The first one is the total variation $\text{TV}(p, q) = \frac{1}{2} \|p - q\|_1$, which can be bounded by Kullback–Leibler divergence $\text{KL}(p||q)$ through Pinsker's inequality

$$\text{TV}(p, q) \leq \sqrt{\frac{1}{2} \text{KL}(p||q)}.$$

The second one is the Wasserstein-2 distance [39]

$$W_2(p, q) := \sqrt{\inf_{v \in \Pi(p, q)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|a - b\|_2^2 dv(a, b)},$$

where $\Pi(p, q)$ denotes the set of the all joint distributions of p and q . The W_2 distance can also be bounded by $\text{KL}(p||q)$ through the Talagrand inequality [40]

$$W_2(p, q) \leq \sqrt{2C_{\text{LSI}} \text{KL}(p||q)},$$

if q satisfies the following log-Sobolev inequality with constant C_{LSI} .

Definition 1 (Log-Sobolev inequality): A probability distribution q satisfies a log-Sobolev inequality with constant C_{LSI} if, for all smooth functions $u : \mathbb{R}^d \rightarrow \mathbb{R}$, it holds that

$$\mathbb{E}_q[u^2 \log u^2] - \mathbb{E}_q[u^2] \log \mathbb{E}_q[u^2] \leq 2C_{\text{LSI}} \mathbb{E}_q[\|\nabla u\|^2].$$

LSI is a powerful tool for sampling convergence analysis when there is no log-concavity. It has nice properties, for instance, it implies the sub-Gaussian concentration property [41]. Moreover, if $s(x)$ is α -strongly-convex, then $\pi(x)$ satisfies the LSI with constant $1/\alpha$ [42]. With the help of the above inequalities, we show that Proposition 1, 2 hold.

Proposition 1: $\text{TV}(\pi, \pi_\beta) \leq \frac{\beta D}{2}$.

Proposition 2: If π satisfies the LSI with constant C_π , then

$$W_2(\pi, \pi_\beta) \leq \sqrt{C_\pi \beta D}.$$

We next compare the smoothing effect of Nesterov smoothing, Moreau envelope and Gaussian smoothing. Define the non-smooth part in our potential as

$$\rho(x) = \max_{y \in \mathcal{Y}} \{ \langle h(x), y \rangle - g(y) \}$$

and denote the Lipschitz constant of $\rho(x)$ as λ_ρ . All three smoothing methods can deal with composite smooth+non-smooth potential, but *Moreau envelope and Gaussian smoothing require the convexity or weakly-convexity of $\rho(x)$, while Nesterov smoothing does not require so.*

a) Moreau envelope: The Moreau-Yosida envelope of a lower semicontinuous (weakly) convex function $\rho(x)$ is defined as

$$\rho_\beta(x) := \min_{y \in \mathbb{R}^d} \left\{ \rho(y) + \frac{1}{2\beta} \|x - y\|_2^2 \right\}.$$

The smoothed composite potential is then written as $s_\beta(x) = f(x) + \rho_\beta(x)$. If ρ is convex, then ρ_β is L_ρ -smooth with $L_\rho = \frac{1}{\beta}$. If ρ is r -weakly convex, and $\beta \leq 1/(2r)$, the same smoothing effect still holds [43], [44]. According to Proposition 1 in [23], there is $\text{TV}(\pi, \pi_\beta) \leq \frac{1}{2}\beta\lambda_\rho^2$, which is implied by $0 \leq s(x) - s_\beta(x) \leq \beta\|\lambda_\rho\|^2/2$.

b) Gaussian smoothing: Unlike the proposed method in [22], where they apply Gaussian smoothing to both the smooth part $f(x)$ and the non-smooth part $\rho(x)$ in $s(x)$, our discussion is based on only smoothing the $\rho_\beta(x)$ here. We find smoothing both parts would introduce a worse dimension dependency in bounding the difference between $\pi(x)$ and $\pi_\beta(x)$. The Gaussian smoothing for $\rho(x)$ is defined as

$$\rho_\beta(x) = \mathbb{E}_\xi [\rho(x + \beta\xi)],$$

where $\xi \sim \mathcal{N}(0, I_{d \times d})$. The smoothing effect of Gaussian smoothing only applies to convex functions. Indeed, according to Lemma 2.2 in [22], if $\rho(x)$ is convex, then $\rho_\beta(x)$ is L_ρ -smooth with $L_\rho = \lambda_\rho d^{1/2}/\beta$. Then by Lemma E.2 in [22], $\text{TV}(\pi, \pi_\beta) \leq \frac{1}{2}\beta\lambda_\rho d^{1/2}$. Additionally, if $f(x)$ is α -strongly convex, by Lemma 3.3 in [22], $W_2(\pi, \pi_\beta) = \tilde{\mathcal{O}}\left(\frac{\beta\lambda_\rho d}{\alpha}\right)$.

Now, if we want to achieve $\text{TV}(\pi, \pi_\beta) = \mathcal{E}/2$, then Moreau envelope results in

$$L_\rho = \frac{\lambda_\rho^2}{\mathcal{E}};$$

Gaussian smoothing needs

$$L_\rho = \frac{\lambda_\rho^2 d}{\mathcal{E}};$$

Nesterov smoothing results in

$$L_\rho = RL_h + \frac{\lambda_h^2 D}{\mathcal{E}\sigma}.$$

Given the same smooth sampling algorithm, the dimension dependency in L_ρ determines the dimension dependency in the final convergence rate. Normally, $\sigma = 1$. Assuming $\lambda_\rho = \lambda_h$, $D = \tilde{\mathcal{O}}(1)$, then Nesterov smoothing effect is the same order as Moreau envelope, and better than Gaussian smoothing. However, Moreau envelope requires more strict assumption on $\rho(x)$, i.e. weakly-convexity; and since the potential is not smooth, solving the proximal operator of Moreau envelope requires additional complexity in each sampling step, even for structured potential in (1).

IV. NON-ASYMPTOTIC CONVERGENCE ANALYSIS

In this section, we present non-asymptotic convergence results in several fundamental scenarios: strongly log-concavity, log-concavity, and log-Sobolev inequality. Denote the underlying distribution of x_k as μ_k . We will give a sufficient number of iterations for μ_k to reach a given error tolerance. In Proposition 3, each complexity result is accompanied by a concrete smoothing algorithm. As discussed in Section II-B, there are multiple algorithm candidates to use for the smooth sampling in Algorithm 1. We will choose the option with the best complexity to our knowledge.

To obtain the non-asymptotic convergence results, we notice that the error, for example, $\text{TV}(\mu_K, \pi)$ can be split into two parts $\text{TV}(\mu_K, \pi_\beta)$ and $\text{TV}(\pi_\beta, \pi)$ by triangular inequality. The former can be bounded using existing convergence results of smooth sampling, such as [36], [35]. Thanks to Proposition 1, 2, the latter is bounded by a carefully chosen β . The convergence in terms of $W_2(\mu_K, \pi)$ follows similarly.

Proposition 3: 1) If f is α -strongly-convex, $\langle h(\cdot), y \rangle$ is convex for any $y \in \mathcal{Y}$, the smooth sampling algorithm in Algorithm 1 is randomized midpoint KLMC [36], and $\beta = \frac{\sqrt{\alpha\mathcal{E}}}{2D}$, then for any $0 < \mathcal{E} < 2\sqrt{\frac{d}{\alpha}}$, the iterate μ_K satisfies $W_2(\mu_K, \pi) \leq \mathcal{E}$ for $K = \tilde{\mathcal{O}}\left(\left(\frac{L_{s_\beta}}{\alpha}\right)^{7/6} \left(\frac{2}{\mathcal{E}}\sqrt{\frac{d}{\alpha}}\right)^{1/3} + \left(\frac{L_{s_\beta}}{\alpha}\right) \left(\frac{2}{\mathcal{E}}\sqrt{\frac{d}{\alpha}}\right)^{2/3}\right)$. Furthermore, if \mathcal{E} is sufficiently small such that $\frac{L_f + RL_h}{\lambda_h^2 D} = \mathcal{O}(1/\mathcal{E})$, we obtain

$$K = \tilde{\mathcal{O}}\left(\frac{\lambda_h^2 D d^{1/3}}{\alpha^{11/6} \mathcal{E}^{5/3}} + \frac{\lambda_h^{7/3} D^{7/6} d^{1/6}}{\alpha^{23/12} \mathcal{E}^{3/2}}\right).$$

2) If f is convex, $\langle h(\cdot), y \rangle$ is convex for any $y \in \mathcal{Y}$, the smooth sampling algorithm is the proximal sampler [35], and $\beta = \frac{\mathcal{E}}{D}$, then for any $0 < \mathcal{E} < \mathcal{O}\left(\frac{\lambda_h^2 D}{L_f + RL_h}\right)$, μ_K satisfies $\text{TV}(\mu_K, \pi) \leq \mathcal{E}$ within iterations

$$K = \mathcal{O}\left(\frac{\lambda_h^2 D \sqrt{d} W_2^2(\mu_0, \pi_\beta)}{\mathcal{E}^3}\right).$$

3) If $\pi(x)$ satisfies LSI with constant C_π , the smooth sampling algorithm is the proximal sampler [35], and $\beta = \frac{\mathcal{E}}{D}$,

then for any $0 < \mathcal{E} < \mathcal{O}\left(\frac{\lambda_h^2 D}{L_f + RL_h}\right)$, μ_K satisfies $\text{TV}(\mu_K, \pi) \leq \mathcal{E}$ within iterations

$$K = \mathcal{O}\left(\frac{\lambda_h^2 DC_\pi \sqrt{d}}{\mathcal{E}}\right).$$

The dimension dependence in the above results matches the corresponding convergence result for smooth sampling. This is because the smoothing parameter β does not depend on the dimension.

V. EXPERIMENTS

In this section, we consider two examples with the special structure (2), which can be rewritten as

$$s(x) = f(x) + \max_{y \in \Delta_{n-1}} \{\langle h(x), y \rangle\},$$

where Δ_{n-1} is the probability simplex on \mathbb{R}^n . To construct the smoothed $s_\beta(x)$, we choose $\|\cdot\|_{\mathcal{Y}}$ as 1-norm $\|\cdot\|_1$, y_0 is the uniform distribution, and $\ell(y) = \log(n) + \sum_{j=1}^n y_j \log y_j$. Thus by [11, Lemma 3], we have $\sigma = 1$ and $D = \log(n)$. We will choose randomized midpoint KLMC as the smooth algorithm.

A. Synthetic example

We consider $s(x)$ to be the summation of a quadratic term and the maximum of absolute values (also see Section 4.4 of [11])

$$s(x) = \|x\|_{\mathcal{X}}^2 + \max_{1 \leq j \leq m} \{|\langle a_j, x \rangle - b_j|\}, \quad a_j \in \mathbb{R}^d, b_j \in \mathbb{R}. \quad (7)$$

It is not difficult to find that $\lambda_p = \max_{1 \leq j \leq m} \|a_j\|_2$ since it is a piece-wise affine function. Denote the matrix A with rows a_j , $b = (b_1, \dots, b_m)$, and $n = 2m$. Then with $\hat{A} = \begin{pmatrix} A \\ -A \end{pmatrix}$ and $\hat{b} = \begin{pmatrix} b \\ -b \end{pmatrix}$, the form (7) is equivalent with

$$s(x) = \|x\|_2^2 + \max_{y \in \Delta_{n-1}} \{\langle \hat{A}x, y \rangle - \langle \hat{b}, y \rangle\}.$$

Clearly, $s(x)$ is strongly-convex, and we have strongly-convex constant $\alpha = 2$, smoothness $L_f = 2, L_h = 0$, and Lipschitz

$$\lambda_h = \|\hat{A}\|_{\mathcal{X}, \mathcal{Y}} = \max_x \{ \max_{1 \leq j \leq m} |\langle a_j, x \rangle| : \|x\|_2 = 1 \} = \max_{1 \leq j \leq m} \|a_j\|_2.$$

Putting these together in Proposition 3, we obtain the total complexity $\tilde{\mathcal{O}}\left(\frac{\lambda_h^{7/3} d^{1/3}}{\mathcal{E}^{5/3}}\right)$, which is experimentally confirmed in Fig. 1.

B. Robust Bayesian logistic regression

We present a real-world example of robust Bayesian logistic regression, where we design a special posterior distribution to make the Bayesian prediction robust to feature perturbation. Given a dataset $\mathcal{L} = \{l_1, \dots, l_S\}$, a model with parameters $x \in \mathbb{R}^n$ and the prior distribution $p_0(x)$, the nominal task is to sample from the posterior distribution

$$\pi_{nom}(x) := p(x|\mathcal{L}) \propto p_0(x)p(\mathcal{L}|x).$$

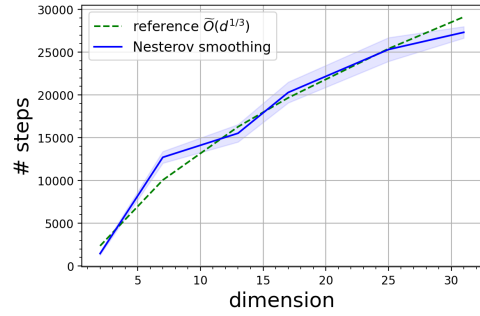


Fig. 1: The number of iterations to reach a fixed error tolerance. The reference dashed line is $2700 d^{1/3} \log(d)$.

Its potential is

$$s_{nom}(x) = -\log p_0(x) - \log p(\mathcal{L}|x).$$

Motivated by *robust approximation problem* [45, Sec. 6.4.2] in optimization, we construct another *worse-case* posterior distribution with the potential

$$s_{wc}(x) = -\log p_0(x) + \max_i \{-\log p(\mathcal{L}_i|x)\},$$

where \mathcal{L}_i is the perturbed version of \mathcal{L} . For example, random noises are added to the features or some features are zeroed. Here we choose to add random Gaussian noise scaled by a noise level to the features, and a larger i corresponds to a greater noise scale. Our setup can be viewed as a Bayesian counterpart of the optimization problem $\min s_{wc}(x)$, which is a *robust approximation problem* aiming to minimize the worst-case objective.

We sample two sets of parameters from distributions π_{nom} and π_{wc} , and test their prediction ability on a series of the perturbed test dataset. They result in two Bayesian predictions/likelihood

$$p_{nom}(l_{test}|\mathcal{L}) = \int p(l_{test}|x)\pi_{nom}(x)dx,$$

$$p_{wc}(l_{test}|\mathcal{L}) = \int p(l_{test}|x)\pi_{wc}(x)dx$$

which are then used to calculate the accuracy. In practice, we use posterior samples to estimate the above integrals/expectations.

The dataset \mathcal{L} we use is "german" dataset from [46], which includes 1000 data and 20 features. The comparison in Fig. 2 illustrates the parameters sampled from π_{wc} give a higher accuracy/log-likelihood when there exist data perturbations.




Fig. 2: Averaged accuracy and the log-likelihood of Bayesian posterior prediction on the test dataset.

VI. CONCLUSION

We consider sampling from a non-smooth distribution $\pi(x) \propto \exp(-s(x))$ that is endowed with the explicit structure (1). Inspired by Nesterov smoothing, we first transform the non-smooth potential $s(x)$ to a smooth potential $s_\beta(x)$ and then apply the smooth sampling algorithms directly to $\pi_\beta(x) \propto \exp(-s_\beta(x))$. Our smoothing technique is compatible with any first-order sampling algorithm for smooth potentials. The error introduced by smoothing can be controlled by a careful choice of smoothing intensity β . We study the non-asymptotic convergence rate under several common conditions and present two examples. The logistic regression example reveals our method can be potentially helpful for robust Bayesian inference. In this paper, we only consider one class of non-smooth potential with max structure (1) for Nesterov smoothing. It is interesting to explore more structures that can benefit from smoothing techniques.

REFERENCES

- [1] A. Golightly and D. J. Wilkinson, "Bayesian sequential inference for nonlinear multivariate diffusions," *Statistics and Computing*, vol. 16, pp. 323–338, 2006.
- [2] A. M. Stuart, "Inverse problems: a Bayesian perspective," *Acta numerica*, vol. 19, pp. 451–559, 2010.
- [3] K. P. Murphy, *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [4] S. Ghosal and A. Van der Vaart, *Fundamentals of nonparametric Bayesian inference*. Cambridge University Press, 2017, vol. 44.
- [5] A. S. Dalalyan, "Theoretical guarantees for approximate sampling from smooth and log-concave densities," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 79, no. 3, pp. 651–676, 2017.
- [6] S. Vempala and A. Wibisono, "Rapid convergence of the unadjusted langevin algorithm: Isoperimetry suffices," *Advances in neural information processing systems*, vol. 32, 2019.
- [7] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust optimization*. Princeton university press, 2009, vol. 28.
- [8] M. Diehl, H. G. Bock, and E. Kostina, "An approximation technique for robust nonlinear optimization," *Mathematical Programming*, vol. 107, no. 1, pp. 213–230, 2006.
- [9] J. Liang and Y. Chen, "A proximal algorithm for sampling from non-smooth potentials," *arXiv preprint arXiv:2110.04597*, 2021.
- [10] S. Chewi, M. A. Erdogdu, M. B. Li, R. Shen, and M. Zhang, "Analysis of langevin monte carlo from poincaré to log-sobolev," *arXiv preprint arXiv:2112.12662*, 2021.
- [11] Y. Nesterov, "Smooth minimization of non-smooth functions," *Mathematical programming*, vol. 103, no. 1, pp. 127–152, 2005.
- [12] A. Wibisono, "Sampling as optimization in the space of measures: The langevin dynamics as a composite optimization problem," in *Conference on Learning Theory*. PMLR, 2018, pp. 2093–3027.
- [13] A. Barp, L. Da Costa, G. França, K. Friston, M. Girolami, M. I. Jordan, and G. A. Pavliotis, "Geometric methods for sampling, optimisation, inference and adaptive agents," *arXiv preprint arXiv:2203.10592*, 2022.
- [14] J. Liang and Y. Chen, "A proximal algorithm for sampling," *arXiv preprint arXiv:2202.13975*, 2022.
- [15] M. A. Erdogdu and R. Hosseinzadeh, "On the convergence of langevin monte carlo: the interplay between tail growth and smoothness," in *Conference on Learning Theory*. PMLR, 2021, pp. 1776–1822.
- [16] D. Nguyen, X. Dang, and Y. Chen, "Unadjusted langevin algorithm for non-convex weakly smooth potentials," *arXiv preprint arXiv:2101.06369*, 2021.
- [17] E. Bernton, "Langevin monte carlo and jko splitting," in *Conference On Learning Theory*. PMLR, 2018, pp. 1777–1798.
- [18] A. Durmus, S. Majewski, and B. Miasojedow, "Analysis of langevin monte carlo via convex optimization," *The Journal of Machine Learning Research*, vol. 20, no. 1, pp. 2666–2711, 2019.
- [19] A. Salim and P. Richtárik, "Primal dual interpretation of the proximal stochastic gradient langevin algorithm," *Advances in Neural Information Processing Systems*, vol. 33, pp. 3786–3796, 2020.
- [20] Y. T. Lee, R. Shen, and K. Tian, "Structured logconcave sampling with a restricted gaussian oracle," in *Conference on Learning Theory*. PMLR, 2021, pp. 2993–3050.
- [21] Y. Chen, S. Chewi, A. Salim, and A. Wibisono, "Improved analysis for a proximal algorithm for sampling," *arXiv preprint arXiv:2202.06386*, 2022.
- [22] N. Chatterji, J. Diakonikolas, M. I. Jordan, and P. Bartlett, "Langevin monte carlo without smoothness," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 1716–1726.
- [23] A. Durmus, E. Moulines, and M. Pereyra, "Efficient bayesian computation by proximal markov chain monte carlo: when langevin meets moreau," *SIAM Journal on Imaging Sciences*, vol. 11, no. 1, pp. 473–506, 2018.
- [24] T. T.-K. Lau and H. Liu, "Bregman proximal langevin monte carlo via bregman-moreau envelopes," in *International Conference on Machine Learning*. PMLR, 2022, pp. 12 049–12 077.
- [25] A. Nemirovski and D. Yudin, "Problem complexity and method efficiency in optimization," 1983.
- [26] Y. Nesterov, *Introductory lectures on convex optimization: A basic course*. Springer Science & Business Media, 2003, vol. 87.
- [27] S. Kakade, S. Shalev-Shwartz, A. Tewari, et al., "On the duality of strong convexity and strong smoothness: Learning applications and matrix regularization," *Unpublished Manuscript*, <http://ttic.uchicago.edu/shai/papers/KakadeShalevTewari09.pdf>, vol. 2, no. 1, p. 35, 2009.
- [28] G. Parisi, "Correlation functions and computer simulations," *Nuclear Physics B*, vol. 180, no. 3, pp. 378–384, 1981.
- [29] X. Cheng, N. S. Chatterji, P. L. Bartlett, and M. I. Jordan, "Underdamped langevin mcmc: A non-asymptotic analysis," in *Conference on learning theory*. PMLR, 2018, pp. 300–323.
- [30] G. O. Roberts and R. L. Tweedie, "Exponential convergence of langevin distributions and their discrete approximations," *Bernoulli*, pp. 341–363, 1996.
- [31] R. M. Neal et al., "Mcmc using hamiltonian dynamics," *Handbook of markov chain monte carlo*, vol. 2, no. 11, p. 2, 2011.
- [32] A. S. Dalalyan and L. Riou-Durand, "On sampling from a log-concave density using kinetic langevin diffusions," *Bernoulli*, vol. 26, no. 3, pp. 1956–1988, 2020.
- [33] Z. Chen and S. S. Vempala, "Optimal convergence rate of hamiltonian monte carlo for strongly logconcave distributions," *Theory of Computing*, vol. 18, no. 1, pp. 1–18, 2022.
- [34] K. Wu, S. Schmidler, and Y. Chen, "Minimax mixing time of the metropolis-adjusted langevin algorithm for log-concave sampling," *arXiv preprint arXiv:2109.13055*, 2021.
- [35] J. Fan, B. Yuan, and Y. Chen, "Improved dimension dependence of a proximal algorithm for sampling," *arXiv preprint arXiv:2302.10081*, 2023.
- [36] R. Shen and Y. T. Lee, "The randomized midpoint method for log-concave sampling," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [37] J. Fan, B. Yuan, J. Liang, and Y. Chen, "Nesterov smoothing for sampling without smoothness," *arXiv preprint arXiv:2208.07459*, 2022.
- [38] S. Bubeck et al., "Convex optimization: Algorithms and complexity," *Foundations and Trends® in Machine Learning*, vol. 8, no. 3-4, pp. 231–357, 2015.
- [39] C. Villani, "Topics in optimal transportation," 2003.
- [40] F. Otto and C. Villani, "Generalization of an inequality by talagrand and links with the logarithmic sobolev inequality," *Journal of Functional Analysis*, vol. 173, no. 2, pp. 361–400, 2000.
- [41] M. Ledoux, "Concentration of measure and logarithmic sobolev inequalities," in *Seminaire de probabilites XXXIII*. Springer, 1999, pp. 120–216.
- [42] D. Bakry and M. Émery, "Diffusions hypercontractives," in *Seminaire de probabilités XIX 1983/84*. Springer, 1985, pp. 177–206.
- [43] A. Böhm and S. J. Wright, "Variable smoothing for weakly convex composite functions," *Journal of optimization theory and applications*, vol. 188, no. 3, pp. 628–649, 2021.
- [44] T. HOHEISEL, M. LABORDE, and A. OBERMAN, "On proximal point-type algorithms for weakly convex functions and their connection to the backward euler method."
- [45] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [46] S. Mika, G. Ratsch, J. Weston, B. Scholkopf, and K.-R. Mullers, "Fisher discriminant analysis with kernels," in *Neural networks for signal processing IX: Proceedings of the 1999 IEEE signal processing society workshop (cat. no. 98th8468)*. Ieee, 1999, pp. 41–48.