



Proximal Oracles for Optimization and Sampling

Jiaming Liang¹  · Yongxin Chen²

Received: 8 July 2025 / Accepted: 10 January 2026

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2026

Abstract

We consider convex optimization with non-smooth objective function and log-concave sampling with non-smooth potential (negative log density). In particular, we study two specific settings where the convex objective/potential function is either Hölder smooth or in hybrid form as the finite sum of Hölder smooth components. To overcome the challenges caused by non-smoothness, our algorithms employ two powerful proximal frameworks in optimization and sampling: the proximal point framework for optimization and the alternating sampling framework (ASF) that uses Gibbs sampling on an augmented distribution. A key component of both optimization and sampling algorithms is the efficient implementation of the proximal map by the regularized cutting-plane method. We establish its iteration-complexity under both Hölder smoothness and hybrid settings using novel convergence analysis, yielding results that are new to the literature. We further propose an adaptive proximal bundle method for non-smooth optimization that employs an aggressive adaptive stepsize strategy, which adjusts stepsizes only when necessary and never rejects iterates. The proposed method is universal since it does not need any problem parameters as input. Additionally, we provide an exact implementation of a proximal sampling oracle, analogous to the proximal map in optimization, along with simple complexity analyses for both the Hölder smooth and hybrid cases, using a novel technique based on a modified Gaussian integral. Finally, we combine this proximal sampling oracle and ASF to obtain a Markov chain Monte Carlo method with non-asymptotic complexity bounds for sampling in Hölder smooth and hybrid settings.

Keywords Non-smooth optimization · Proximal point method · Universal method · High-dimensional sampling · Markov chain Monte Carlo · Complexity analysis

Communicated by Shoham Sabach.

✉ Jiaming Liang
jiaming.liang@rochester.edu

Yongxin Chen
yongchen@gatech.edu

¹ Goergen Institute for Data Science and Artificial Intelligence & Department of Computer Science, University of Rochester, Rochester, NY 14620, USA

² School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

1 Introduction

We are interested in convex optimization problems

$$\min_{x \in \mathbb{R}^d} f(x) \quad (1)$$

as well as log-concave sampling problems

$$\text{sample } \nu(x) \propto \exp(-f(x)), \quad (2)$$

where $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is convex but not necessarily smooth. In sampling, a potential of the distribution $\nu(x)$ is defined as the negative log-density, which is $f(x)$ up to a constant.

Optimization and sampling are two of the most important algorithmic tools at the interface of data science and computation. Optimization has been extensively studied across a wide range of fields, including machine learning, communications, and supply chain management. Over the past two decades, particular attention has been devoted to gradient-based first-order methods. Many classical ideas have been revisited and extended to large-scale optimization, such as the randomized coordinate descent method [52], the primal–dual hybrid gradient method [5], and the extragradient method [30]. Drawing samples from a given (often unnormalized) probability density plays a crucial role in many scientific and engineering problems that face uncertainty (either physically or algorithmically). Sampling algorithms are widely used in many areas such as statistical inference/estimation, operations research, physics, biology, and machine learning, etc [2, 11, 12, 16, 25, 26, 31, 63]. For instance, in Bayesian inference, one draws samples from the posterior distribution to infer its mean, covariance, or other important statistics. Sampling is also heavily used in molecular dynamics to discover new molecular structures.

This work is along the recent line of research that lies in the interface of sampling and optimization [10, 61]. Indeed, sampling is closely related to optimization. On the one hand, optimization can be viewed as the limiting case of sampling from the distribution $\exp(-f(x)/T)$ as the temperature parameter T (which represents the level of randomness) approaches zero. In this limit, the probability mass increasingly concentrates around the minimizers of $f(x)$. On the other hand, sampling $\nu(x)$ has an optimization interpretation [24, 66, 68]: the Langevin dynamics in space corresponds to the Fokker-Planck equation, which is the gradient flow of the relative entropy functional (with respect to ν) in the space of measures with the Wasserstein metric. The popular gradient-based Markov chain Monte Carlo (MCMC) methods such as Langevin Monte Carlo (LMC) [7, 20, 55, 57], Metropolis-adjusted Langevin algorithm (MALA) [3, 56, 57], and Hamiltonian Monte Carlo (HMC) [49] resemble the gradient-based algorithms in optimization and can be viewed as the sampling counterparts of them.

The goal of this paper is to develop efficient proximal algorithms to solve optimization problems (1) as well as to draw samples from potentials (2), where both f in (1) and (2) lack smoothness (i.e., when f does not have Lipschitz continuous gradient). In particular, we consider two settings where the convex objective/potential function f is either Hölder smooth (i.e., the (sub)gradient f' is Hölder-continuous with exponent $\alpha \in [0, 1]$) or a hybrid function with multiple Hölder smooth components. The core of both proximal optimization and sampling algorithms lies in the proximal map of f . We first develop a generic and efficient implementation of this proximal map. Building on it, we design an adaptive proximal bundle method to solve problem (1). Furthermore, by combining the proximal map of f with rejection sampling, we propose a highly efficient approach to realize a proximal sampling oracle, which is used in a proximal sampling framework [6, 33] in the same spirit as the proximal point method for optimization. With those proximal oracles for optimization and sampling in hand, we are finally able to establish the complexity to sample from densities with non-smooth potentials.

We summarize our contributions as follows.

- i) We analyze the complexity bounds for implementing the proximal map of f using the regularized cutting-plane method in both Hölder smooth and hybrid settings (Sect. 3). The complexity analyses for both Hölder smooth and hybrid cases, presented in Subsections 3.1 and 3.2, respectively, are novel contributions to the literature and employ proof techniques distinct from existing works such as [9, 27, 41, 42, 69].
- ii) We develop an adaptive proximal bundle method (APBM) using the regularized cutting-plane method and a novel adaptive stepsize strategy in the proximal point method, and establish the complexity bound for Hölder smooth optimization (Sect. 4). APBM is a universal method as it does not need any problem-dependent parameters as input. In contrast to standard universal methods based on conservative line searches on stepsizes, such as the universal primal gradient method of [53], APBM has the benefit of adjusting stepsizes only when necessary and never rejects iterates.
- iii) We propose an efficient scheme to realize the proximal sampling oracle that lacks smoothness and establish novel techniques to bound its complexity. Combining the proximal sampling oracle and the proximal sampling framework, we obtain a general proximal sampling algorithm for convex Hölder smooth and hybrid potentials. Finally, we establish complexity bounds for the proximal sampling algorithm in both cases (Sect. 5). The complexity bounds presented in Sect. 5 are similar to those in [13]; however, they are derived under the assumption of an exact proximal sampling oracle, whereas [13] considers an inexact implementation of the oracle. The contributions of Sect. 5 lie in providing much simpler complexity analyses for the exact realization of the proximal sampling oracle in both the Hölder smooth and hybrid cases, compared to the existing analyses in [37, 38].

It is worth noting that this paper does not aim to establish the optimal complexity of universal methods or to improve the complexity of proximal sampling algorithms. Instead, it develops a regularized cutting-plane method as an efficient implementation of the proximal oracle used in both proximal optimization and sampling, and

demonstrates its interesting applications in universal methods and proximal sampling algorithms.

2 Proximal Optimization and Sampling

The proximal point framework (PPF), proposed in [44] and further developed in [58, 59] (see [54] for a modern and comprehensive monograph), is a general class of optimization algorithms that involve solving a sequence of subproblems of the form

$$x_{k+1} \leftarrow \operatorname{argmin} \left\{ f(x) + \frac{1}{2\eta} \|x - x_k\|^2 : x \in \mathbb{R}^d \right\}, \quad (3)$$

where $\eta > 0$ is a prox stepsize and \leftarrow means the subproblem can be solved either exactly or approximately. When the exact solution is available, we denote

$$x_{k+1} = \operatorname{prox}_{\eta f}(x_k),$$

where $\operatorname{prox}_f(\cdot)$ is called a proximal map of f and defined as

$$\operatorname{prox}_f(y) := \operatorname{argmin} \left\{ f(x) + \frac{1}{2} \|x - y\|^2 : x \in \mathbb{R}^d \right\}. \quad (4)$$

If the subproblem (3) does not admit a closed-form solution, it can usually be solved with standard or specialized iterative methods.

Many classical first-order methods in optimization, such as the proximal gradient method, the proximal subgradient method, the primal-dual hybrid gradient method of [5] (also known as the Chambolle-Pock method), the extra gradient method of [30] are instances of PPF. It is worth noting that, by showing that the alternating direction method of multipliers (ADMM) as an instance of PPF, [47] gives the first iteration-complexity result of ADMM for solving a class of linearly constrained convex programming problems.

Another example of PPF is the proximal bundle method, which was first proposed in [34, 35, 45, 67] and further developed in [8, 9, 14, 27, 41, 42, 60, 64, 69]. Notably, inspired by the PPF viewpoint, papers [41, 42] develop a variant of the proximal bundle method and establish the optimal iteration-complexity, which is the first optimal complexity result for proximal bundle methods.

Recent works [28, 29, 40, 43] have also applied PPF to solve weakly convex optimization and weakly convex-concave min-max problems.

Proximal map in sampling. Sampling shares many similarities with optimization. An interesting connection between the two problems is through the algorithm design and analysis from the perspective of PPF. The alternating sampling framework (ASF) introduced in [33] is a generic framework for sampling from a distribution $\pi^X(x) \propto \exp(-f(x))$. Analogous to PPF in optimization, ASF with stepsize $\eta > 0$ repeats the two steps as in Algorithm 1.

Algorithm 1 Alternating Sampling Framework [33]

1. Sample $y_k \sim \pi^{Y|X}(y | x_k) \propto \exp\left(-\frac{1}{2\eta}\|x_k - y\|^2\right)$
2. Sample $x_{k+1} \sim \pi^{X|Y}(x | y_k) \propto \exp\left(-f(x) - \frac{1}{2\eta}\|x - y_k\|^2\right)$

ASF is a special case of Gibbs sampling [17] of the joint distribution

$$\pi(x, y) \propto \exp\left(-f(x) - \frac{1}{2\eta}\|x - y\|^2\right).$$

Starting from the original paper [33] that proposes ASF, subsequent works have refined and extended this framework. In particular, [6] provides an improved theoretical analysis of ASF, and [70] studies Gibbs sampling based on ASF for structured log-concave distributions over networks. In Algorithm 1, sampling y_k given x_k in step 1 can be easily done since $\pi^{Y|X}(y | x_k) = \mathcal{N}(x_k, \eta I)$ is a simple Gaussian distribution. Sampling x_{k+1} given y_k in step 2 is however a nontrivial task; it corresponds to the so-called restricted Gaussian oracle (RGO) for f introduced in [33], which is defined as follows.

Definition 2.1 Given a point $y \in \mathbb{R}^d$ and stepsize $\eta > 0$, the RGO for $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is a sampling oracle that returns a random sample from a distribution proportional to $\exp(-f(\cdot) - \|\cdot - y\|^2/(2\eta))$.

RGO is an analog of the proximal map (4) in optimization. To use ASF in practice, one needs to efficiently implement RGO. Some examples of f that admit a computationally efficient RGO have been presented in [48, 62]. These instances of f have simple structures such as coordinate-separable regularizers, ℓ_1 -norm, and group Lasso. To apply ASF on a general potential function f , developing an efficient implementation of the RGO is essential.

A rejection sampling-based implementation of RGO for a general convex nonsmooth potential function f with bounded Lipschitz constant is given in [37].

If the stepsize η is small enough, then it only takes a constant number of rejection steps to generate a sample according to RGO in expectation. Another exact realization of RGO is provided in [38] for nonconvex hybrid potential f satisfying Hölder continuous conditions. It is also shown that the expected number of rejections to implement RGO is a small constant if η is small enough. Other inexact realizations of RGO based on approximate rejection sampling are studied in [13, 18]. See Table 1 for a clear comparison. In all these implementations, a key step is realizing the proximal map (4). It is worth noting that [38] also connects ASF with other well-known Langevin-type sampling algorithms such as Langevin Monte Carlo (LMC) and Proximal Langevin Monte Carlo (PLMC) via RGO. In a nutshell, [38] shows that both LMC and PLC are instances of ASF but with approximate implementations of RGO, which always accept the sample from the proposal distribution without rejection. Hence, this provides an alternative interpretation of why the samples generated by LMC are biased, while those produced by ASF are unbiased.

Based on the cutting-plane method, this paper develops a generic and efficient implementation of the proximal map (4) and applies the proximal map in both opti-

Table 1 Comparison of different RGO implementations and corresponding stepsizes

Papers	RGO implementation	Stepsize η
[37, 38]	Exact	Small
[13, 18]	Approximate	Large

mization and sampling. For optimization, we use this proximal map and an adaptive stepsize rule to design a universal bundle method. For sampling, we combine this proximal map and rejection sampling to realize the RGO, and then propose a practical and efficient proximal sampling algorithm based on it.

For both optimization and sampling, we consider two specific scenarios: (1) f is Hölder smooth, i.e., f satisfies

$$\|f'(u) - f'(v)\| \leq L_\alpha \|u - v\|^\alpha, \quad \forall u, v \in \mathbb{R}^d, \quad (5)$$

where f' denotes a subgradient of f , $\alpha \in [0, 1]$, and $L_\alpha > 0$; and (2) f is a hybrid function of Hölder smooth components, i.e., f satisfies

$$\|f'(u) - f'(v)\| \leq \sum_{i=1}^n L_{\alpha_i} \|u - v\|^{\alpha_i}, \quad \forall u, v \in \mathbb{R}^d, \quad (6)$$

where $\alpha_i \in [0, 1]$ and $L_{\alpha_i} > 0$ for every $1 \leq i \leq n$. When $\alpha = 0$, (5) reduces to a Lipschitz continuous condition, and when $\alpha = 1$, it reduces to a smoothness condition. It follows from (5) and (6) that for every $u, v \in \mathbb{R}^d$,

$$f(u) - f(v) - \langle f'(v), u - v \rangle \leq \frac{L_\alpha}{\alpha + 1} \|u - v\|^{\alpha+1}, \quad (7)$$

and

$$f(u) - f(v) - \langle f'(v), u - v \rangle \leq \sum_{i=1}^n \frac{L_{\alpha_i}}{\alpha_i + 1} \|u - v\|^{\alpha_i+1}. \quad (8)$$

The proof is given in Appendix A.

Example. Consider the ℓ_p regression problem with data $\{(a_i, b_i)\}_{i=1}^n$ where $a_i \in \mathbb{R}^d$ and $b_i \in \mathbb{R}$ for $i = 1, \dots, n$,

$$f(x) = \frac{1}{n} \sum_{i=1}^n |a_i^\top x - b_i|^p, \quad 1 \leq p \leq 2. \quad (9)$$

Define $\phi(t) = |t|^p$, then $\phi'(t) = p \operatorname{sign}(t) |t|^{p-1}$ and

$$f'(x) = \frac{1}{n} \sum_{i=1}^n \phi'(a_i^\top x - b_i) a_i.$$

It is shown in Lemma A.4 of Appendix A that ϕ' is Hölder continuous with exponent $p - 1$ and constant $p2^{2-p}$. For any $x, y \in \mathbb{R}^d$, let $u_i = a_i^\top x - b_i$ and $v_i = a_i^\top y - b_i$, using the Hölder continuity of ϕ' , we derive

$$\begin{aligned} \|f'(x) - f'(y)\| &= \left\| \frac{1}{n} \sum_{i=1}^n (\phi'(u_i) - \phi'(v_i)) a_i \right\| \\ &\leq \frac{1}{n} \sum_{i=1}^n |\phi'(u_i) - \phi'(v_i)| \|a_i\| \leq \frac{p2^{2-p}}{n} \left(\sum_{i=1}^n \|a_i\|^p \right) \|x - y\|^{p-1}. \end{aligned}$$

Hence, f satisfies the Hölder smoothness condition (5) with

$$\alpha = p - 1, \quad L_\alpha = \frac{p2^{2-p}}{n} \sum_{i=1}^n \|a_i\|^p.$$

The ℓ_p regression can be extended to mixed-exponent regression as an example of the hybrid case (6), where

$$f(x) = \frac{1}{n} \sum_{i=1}^n |a_i^\top x - b_i|^{p_i}, \quad 1 \leq p_i \leq 2, \quad (10)$$

and

$$\alpha_i = p_i - 1, \quad L_{\alpha_i} = \frac{p_i 2^{2-p_i}}{n} \|a_i\|^{p_i}.$$

The above objective functions f in (9) and (10) can also appear as the potential energy in Bayesian inference. Instead of minimizing $f(x)$ to obtain a point estimate (e.g., the maximum a posteriori or MAP solution), one may consider sampling $v(x) \propto \exp(-f(x))$ for quantifying uncertainty around the MAP solution.

Throughout the analysis in this paper, we use the following notation. When presenting complexity results, $\mathcal{O}(\cdot)$ denotes the standard ‘‘big-O’’ notation, while $\tilde{\mathcal{O}}(\cdot)$ suppresses polylogarithmic factors. We also write $a \asymp b$ to indicate that a and b are of the same order, i.e., there exist positive constants $c_1, c_2 > 0$ such that $c_1 a \leq b \leq c_2 a$.

3 Algorithm and Complexities for the Proximal Subproblem

The proximal subproblem (3) generally does not admit a closed-form solution. We design an iterative method that approximately solves (3) and derive the corresponding iteration-complexities for Hölder smooth and hybrid f in Subsections 3.1 and 3.2, respectively.

Given a point $y \in \mathbb{R}^d$, we consider the optimization problem

$$f_y^\eta(x^*) = \min \left\{ f_y^\eta(x) = f(x) + \frac{1}{2\eta} \|x - y\|^2 : x \in \mathbb{R}^d \right\} \quad (11)$$

and aim at obtaining a δ -solution, i.e., a point \bar{x} such that $f_y^\eta(\bar{x}) - f_y^\eta(x^*) \leq \delta$. In both Hölder smooth and hybrid settings, we use a regularized cutting-plane method (Algorithm 2), which is usually used in the proximal bundle method [41, 42] for solving convex non-smooth optimization problems.

We remark that though Algorithm 2 is widely used in the proximal bundle method and is not new, the complexity analyses (i.e., Theorems 3.1 and 3.2) for Hölder smooth and hybrid functions f are lacking.

Since the prox center y is fixed throughout this section, we simplify the notation by writing f_y^η as f^η in this section to ease readability.

Algorithm 2 Regularized Cutting-plane Method

Require: Let $y \in \mathbb{R}^d$, $\eta > 0$, and $\delta > 0$ be given, and set $x_0 = \tilde{x}_0 = y$, $j = 1$, and $f_0^\eta(x_0) = -\infty$.

while $f^\eta(\tilde{x}_{j-1}) - f_{j-1}^\eta(x_{j-1}) > \delta$ **do**

$$f_j(x) = \max \{f(x_i) + \langle f'(x_i), x - x_i \rangle : 0 \leq i \leq j-1\}, \quad (12)$$

$$x_j = \operatorname{argmin} \left\{ f_j^\eta(x) := f_j(x) + \frac{1}{2\eta} \|x - y\|^2 : x \in \mathbb{R}^d \right\}, \quad (13)$$

$$\tilde{x}_j = \operatorname{argmin} \{f^\eta(x) : x \in \{x_j, \tilde{x}_{j-1}\}\}, \quad (14)$$

$$j \leftarrow j + 1.$$

end while

return $J = j - 1$, x_J , and \tilde{x}_J .

The basic idea of Algorithm 2 is to approximate f with piece-wise affine functions constructed by a collection of cutting-planes and solve the resulting simplified problem (13). As the approximation becomes more and more accurate, the best approximate solution \tilde{x}_j converges to the solution x^* to (11).

Subproblem (13) can be reformulated into convex quadratic programming with j affine constraints and hence is solvable.

The following technical lemma summarizes basic properties of Algorithm 2. It is useful in the complexity analysis for both optimization and sampling.

Lemma 3.1 Assume f is convex.

For every $j \geq 1$, define

$$\delta_j := f^\eta(\tilde{x}_j) - f_j^\eta(x_j). \quad (15)$$

Let J , x_J , \tilde{x}_J be the outputs of Algorithm 2, then the following statements hold:

- $\{f_j\}$ serves as a sequence of non-decreasing lower approximations of f : $f_j(x) \leq f_{j+1}(x)$ and $f_j(x) \leq f(x)$, $\forall x \in \mathbb{R}^d$ and $\forall j \geq 1$;
- direct consequence of (13): $f_j^\eta(x_j) + \|x - x_j\|^2/(2\eta) \leq f_j^\eta(x)$, $\forall x \in \mathbb{R}^d$ and $\forall j \geq 1$;
- $\{\delta_j\}$ is a decreasing sequence: $\delta_J \leq \delta$ and $\delta_{j+1} + \frac{1}{2\eta} \|x_{j+1} - x_j\|^2 \leq \delta_j$, $\forall j \geq 1$;
- solution guarantee for x_J and \tilde{x}_J : $f^\eta(\tilde{x}_J) - f^\eta(x) \leq \delta - \frac{1}{2\eta} \|x_J - x\|^2$, $\forall x \in \mathbb{R}^d$;

e) optimality condition of (11): $-\frac{1}{\eta}(x^* - y) \in \partial f(x^*)$ where ∂f denotes the subdifferential of f .

Proof a) The first inequality follows from the definition of f_j in step 2 of Algorithm 2. The second inequality directly follows from the definition of f_j and the convexity of f .

b) Noting that f_j^η as the objective function of (13) is $(1/\eta)$ -strongly convex, it thus follows from Theorem 5.25 of [1] that

$$f_j^\eta(x) - f_j^\eta(x_j) \geq \frac{1}{2\eta} \|x - x_j\|^2, \quad \forall x \in \mathbb{R}^d.$$

Hence, this statement follows.

c) This first inequality immediately follows from (15) and step 4 of Algorithm 2. Using the first inequality in 3.1(a) and 3.1(b) with $x = x_{j+1}$, we obtain

$$f_{j+1}^\eta(x_{j+1}) \geq f_j^\eta(x_{j+1}) \geq f_j^\eta(x_j) + \frac{1}{2\eta} \|x_{j+1} - x_j\|^2.$$

This inequality, the definition of \tilde{x}_j in (14), and the definition of δ_j in (15) imply that

$$\begin{aligned} \delta_{j+1} &= f^\eta(\tilde{x}_{j+1}) - f_{j+1}^\eta(x_{j+1}) \leq f^\eta(\tilde{x}_j) - f_j^\eta(x_j) - \frac{1}{2\eta} \|x_{j+1} - x_j\|^2 \\ &= \delta_j - \frac{1}{2\eta} \|x_{j+1} - x_j\|^2. \end{aligned}$$

d) Using the second inequality in (a), (b) with $j = J$, and the first inequality in (c), we have

$$\begin{aligned} f(\tilde{x}_J) - f(x) + \frac{1}{2\eta} \|x - x_J\|^2 &\stackrel{(a)}{\leq} f(\tilde{x}_J) - f_J(x) + \frac{1}{2\eta} \|x - x_J\|^2 \\ &\stackrel{(b)}{\leq} f(\tilde{x}_J) - f_J^\eta(x_J) + \frac{1}{2\eta} \|x - y\|^2 \stackrel{(c)}{\leq} \delta - \frac{1}{2\eta} \|\tilde{x}_J - y\|^2 + \frac{1}{2\eta} \|x - y\|^2. \end{aligned}$$

This statement then follows from rearranging the terms and the definition of f^η in (11).

e) This statement directly follows from the first-order optimality condition of (11). \square

Clearly, when Algorithm 2 terminates, the output \tilde{x}_J is a δ -solution to (11). To see this, note that, using the first inequality in Lemma 3.1(c), (13), and the fact that $f_J^\eta(\cdot) \leq f^\eta(\cdot)$, we have

$$f^\eta(\tilde{x}_J) \leq \delta + f_J^\eta(x_J) \stackrel{(13)}{\leq} \delta + f_J^\eta(x^*) \leq \delta + f^\eta(x^*).$$

It is also easy to see that δ_j is computable upper bound on the gap $f^\eta(\tilde{x}_j) - f^\eta(x^*)$. Hence, Algorithm 2 terminates when $\delta_j \leq \delta$.

3.1 Complexity for Hölder Smooth Optimization

This subsection is devoted to the complexity analysis of Algorithm 2 for solving (11) where f is Hölder smooth, i.e., satisfying (5). The following lemma provides basic recursive formulas and is the starting point of the analysis of Algorithm 2.

Lemma 3.2 *Assume f is convex and L_α -Hölder smooth. Then, for every $j \geq 1$, the following statements hold:*

- a) $\delta_j \leq \frac{L_\alpha}{\alpha+1} \|x_j - x_{j-1}\|^{\alpha+1}$;
- b) $\delta_{j+1} + \frac{1}{2\eta} \left(\frac{\alpha+1}{L_\alpha} \delta_{j+1} \right)^{\frac{2}{\alpha+1}} \leq \delta_j$.

Proof a) It follows from the definition of δ_j in (15) and the definition of \tilde{x}_j in (14) that

$$\begin{aligned} \delta_j &\stackrel{(15)}{=} f^\eta(\tilde{x}_j) - f_j^\eta(x_j) \stackrel{(14)}{\leq} f^\eta(x_j) - f_j^\eta(x_j) = f(x_j) - f_j(x_j) \\ &\leq f(x_j) - f(x_{j-1}) - \langle f'(x_{j-1}), x_j - x_{j-1} \rangle \\ &\leq \frac{L_\alpha}{\alpha+1} \|x_j - x_{j-1}\|^{\alpha+1}, \end{aligned}$$

where the second inequality is due to the definition of f_j in the step 2 of Algorithm 2, and the third inequality is due to (7) with $(u, v) = (x_j, x_{j-1})$.

b) This statement directly follows from a) and the second inequality in Lemma 3.1(c). \square

We know from Lemma 3.1(c) that $\{\delta_j\}_{j \geq 1}$ is non-increasing. The next proposition gives a bound on j so that $\delta_j \leq \delta$, i.e., the termination criterion in step 4 of Algorithm 2 is satisfied.

Proposition 3.1 *Define*

$$\beta := \frac{1}{2\eta} \left(\frac{\alpha+1}{L_\alpha} \right)^{\frac{2}{\alpha+1}} \delta^{\frac{1-\alpha}{\alpha+1}}, \quad j_0 = 1 + \left\lceil \frac{1+\beta}{\beta} \log \left(\frac{\delta_1}{\delta} \right) \right\rceil. \quad (16)$$

Then, the following statements hold:

- a) if $\delta_j > \delta$, then $(1+\beta)\delta_j \leq \delta_{j-1}$;
- b) $\delta_j \leq \delta$ for every $j \geq j_0$.

As a consequence, the iteration count J in Algorithm 2 satisfies $J \leq j_0$.

Proof a) Using the definition of β in (16), the assumption that $\delta_j > \delta$, and Lemma 3.2(b), we obtain

$$(1+\beta)\delta_j = \delta_j + \frac{1}{2\eta} \left(\frac{\alpha+1}{L_\alpha} \right)^{\frac{2}{\alpha+1}} \delta^{\frac{1-\alpha}{\alpha+1}} \delta_j \leq \delta_j + \frac{1}{2\eta} \left(\frac{\alpha+1}{L_\alpha} \delta_j \right)^{\frac{2}{\alpha+1}} \leq \delta_{j-1}.$$

b) Since $\{\delta_j\}_{j \geq 1}$ is non-increasing, it suffices to prove that $\delta_{j_0} \leq \delta$. We prove this statement by contradiction. Suppose that $\delta_{j_0} > \delta$, then we have $\delta_j > \delta$ for $j \leq j_0$. Hence, statement (a) holds for $j \leq j_0$. Using this conclusion repeatedly and the fact that $\tau \leq \exp(\tau - 1)$ with $\tau = 1/(1 + \beta)$, we have

$$\delta_{j_0} \leq \frac{1}{(1 + \beta)^{j_0-1}} \delta_1 \leq \exp\left(-\frac{\beta}{1 + \beta}(j_0 - 1)\right) \delta_1 \leq \delta,$$

where the last inequality is due to the definition of j_0 in (16). This contradicts with the assumption that $\delta_{j_0} > \delta$, and hence we prove this statement. \square

The following result shows that δ_1 is bounded from above, and hence the bound in Proposition 3.1 is meaningful.

Lemma 3.3 *For a given $y \in \mathbb{R}^d$, we have*

$$\delta_1 \leq \frac{L_\alpha \eta^{\alpha+1}}{\alpha + 1} \|f'(y)\|^{\alpha+1}.$$

Proof Following the optimality condition of (13) with $j = 1$, we have $x_0 - x_1 = \eta f'(x_0) = \eta f'(y)$. This identity and Lemma 3.2(a) with $j = 1$ then imply that the lemma holds. \square

We now conclude the iteration-complexity bound for Algorithm 2.

Theorem 3.1 *Algorithm 2 takes $\tilde{\mathcal{O}}\left(\eta L_\alpha^{\frac{2}{\alpha+1}} \left(\frac{1}{\delta}\right)^{\frac{1-\alpha}{\alpha+1}} + 1\right)$ iterations to terminate.*

Proof This theorem follows directly from Proposition 3.1 and Lemma 3.3. \square

3.2 Complexity for Hybrid Optimization

This subsection is devoted to the complexity analysis of Algorithm 2 for solving (11) where f is a hybrid function satisfying (6). The following lemma is an analogue of Lemma 3.2 and provides key recursive formulas for δ_j , which is defined in (15).

Lemma 3.4 *Assume f is convex and satisfies (6). For $\delta > 0$, define*

$$M = \sum_{i=1}^n \frac{L_{\alpha_i}^{\frac{2}{\alpha_i+1}}}{[(\alpha_i + 1)\delta]^{\frac{1-\alpha_i}{\alpha_i+1}}}. \quad (17)$$

Then, for every $j \geq 1$, the following statements hold:

- a) $\delta_j \leq \frac{M}{2} \|x_j - x_{j-1}\|^2 + \sum_{i=1}^n (1 - \alpha_i) \frac{\delta}{2}$;
- b) $\left(1 + \frac{1}{\eta M}\right) (\delta_{j+1} - \sum_{i=1}^n (1 - \alpha_i) \frac{\delta}{2}) \leq \delta_j - \sum_{i=1}^n (1 - \alpha_i) \frac{\delta}{2}$.

Proof a) Following a similar argument as in the proof of Lemma 3.2(a) with (7) replaced by (8), we have

$$\delta_j \leq \sum_{i=1}^n \frac{L_{\alpha_i}}{\alpha_i + 1} \|u - v\|^{\alpha_i + 1}. \quad (18)$$

Using the Young's inequality $ab \leq a^p/p + b^q/q$ with

$$a = \frac{L_\alpha}{(\alpha + 1)\delta^{\frac{1-\alpha}{2}}} \|x_j - x_{j-1}\|^{\alpha+1}, \quad b = \delta^{\frac{1-\alpha}{2}}, \quad p = \frac{2}{\alpha + 1}, \quad q = \frac{2}{1 - \alpha},$$

we obtain

$$\frac{L_\alpha}{\alpha + 1} \|x_j - x_{j-1}\|^{\alpha+1} \leq \frac{L_\alpha^{\frac{2}{\alpha+1}}}{2[(\alpha + 1)\delta]^{\frac{1-\alpha}{\alpha+1}}} \|x_j - x_{j-1}\|^2 + \frac{(1 - \alpha)\delta}{2}.$$

Combining the above inequality and (18), and using the definition of M in (17), we prove the statement.

b) It immediately follows from (a) and the second inequality in Lemma 3.1(c) that

$$\delta_{j+1} + \frac{1}{\eta M} \left(\delta_{j+1} - \sum_{i=1}^n (1 - \alpha_i) \frac{\delta}{2} \right) \leq \delta_{j+1} + \frac{1}{2\eta} \|x_{j+1} - x_j\|^2 \leq \delta_j,$$

and hence the statement follows. \square

The following lemma gives an upper bound on δ_1 similar to Lemma 3.3.

Lemma 3.5 *For a given $y \in \mathbb{R}^d$, we have*

$$\delta_1 \leq \sum_{i=1}^n \frac{L_{\alpha_i}}{\alpha_i + 1} \|f'(y)\|^{\alpha_i + 1}.$$

Proof This lemma follows from a similar argument as in the proof of Lemma 3.3. \square

The following proposition is the key result in establishing the iteration-complexity of Algorithm 2.

Proposition 3.2 *We have $\delta_j \leq \delta$, for every j such that*

$$j \geq (1 + \eta M) \log \left(\frac{2\delta_1}{\delta} \right). \quad (19)$$

Proof Let

$$\tau = \frac{\eta M}{1 + \eta M}, \quad (20)$$

then Lemma 3.4(b) becomes

$$\delta_{j+1} - \sum_{i=1}^n (1 - \alpha_i) \frac{\delta}{2} \leq \tau \left(t_j - \sum_{i=1}^n (1 - \alpha_i) \frac{\delta}{2} \right).$$

Using the above inequality repeatedly and the fact that $\tau \leq \exp(\tau - 1)$, we have for every $j \geq 1$,

$$\delta_j - \frac{(1 - \alpha)\delta}{2} \leq \tau^{j-1} \left(\delta_1 - \frac{(1 - \alpha)\delta}{2} \right) \leq \tau^{j-1} \delta_1 \leq \exp\{(\tau - 1)(j - 1)\} \delta_1.$$

Hence, it is easy to see that $\delta_j \leq \delta$ if $j \geq \frac{1}{1-\tau} \log \left(\frac{2\delta_1}{\delta} \right)$. Using the definition of τ in (20), we have if j is as in (19), then $\delta_j \leq \delta$. \square

We are ready to present the complexity bound for Algorithm 2.

Theorem 3.2 *Algorithm 2 takes $\tilde{\mathcal{O}}(\eta M + 1)$ iterations to terminate, where M is as in (17).*

Proof This theorem follows directly from Proposition 3.2 and Lemma 3.5. \square

3.3 Implementation of Algorithm 2

This subsection presents the simulation results of Algorithm 2 on solving the regularized subproblem (11) for two objective functions f : quadratic programming (QP) and ℓ_p regression. In both cases, the subgradient f' is computed by automatic differentiation via Zygote.jl [22], and the subproblem (13) is reformulated as a QP and solved using Clarabel.jl [19]. Numerical simulations are conducted on an i9-13900k desktop with 64 GB of RAM

Quadratic Programming We first consider the unconstrained QP problem

$$f(x) = \frac{1}{2} x^\top Q x + \langle c, x \rangle$$

where $Q \in \mathbb{S}_+^d$ and $c \in \mathbb{R}^d$. We generate $Q = AA^\top / \|AA^\top\|_\infty$ where $A \in \mathbb{R}^{d \times d}$ has normally distributed entries and $\|AA^\top\|_\infty = \max_{ij} |(AA^\top)_{ij}|$ is the entrywise infinity norm. The linear term c and point y are also entrywise normally distributed. The dimension d is set to be 1000.

Note that (11) for QP has the closed-form solution

$$\underset{x \in \mathbb{R}^n}{\operatorname{argmin}} \left\{ \frac{1}{2} x^\top Q x + \langle c, x \rangle + \frac{1}{2\eta} \|x - y\|^2 \right\} = (Q + \eta^{-1} I)^{-1} (\eta^{-1} y - c),$$

hence we can compare the progress of Algorithm 2 against the true minimum. We run Algorithm 2 until the condition $\delta_j < 10^{-6}$ is satisfied. Figure 1 shows the function value decrease of the minimum value iterate $f_y^\eta(\tilde{x}_j)$ versus the optimal value $f_y^\eta(x^*)$

Fig. 1 Proximal subproblem progress of Algorithm 2 in quadratic programming

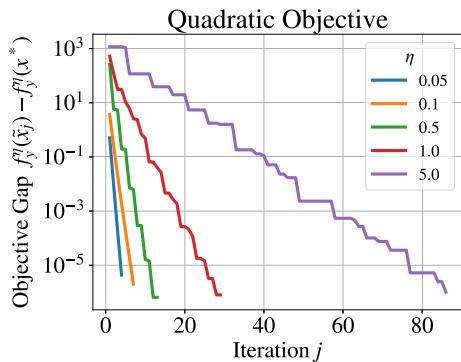
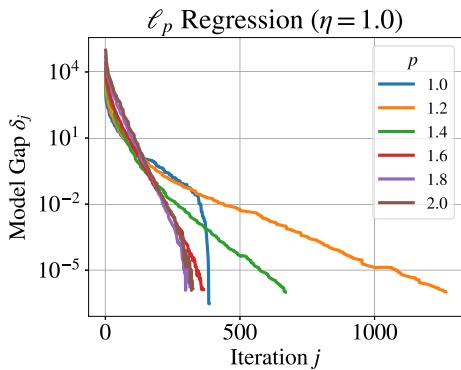


Fig. 2 Proximal subproblem progress of Algorithm 2 in ℓ_p regression



with varying η . Noting that Algorithm 2 requires more iterations as η increases, this observation is consistent with the complexity bound (proportional to η) stated in Theorem 3.1.

ℓ_p Regression We next consider ℓ_p regression, where the objective f is of the form

$$f(x) = \|Ax - b\|_p^p,$$

where $A \in \mathbb{R}^{n \times d}$ and $b \in \mathbb{R}^n$ have normally distributed entries and A is again divided by its entrywise infinity norm. We set $d = 100$ and $n = 500$ for testing. The point $y \in \mathbb{R}^d$ is entrywise normally distributed, and is identical for all p values tested. Algorithm 2 is terminated when $\delta_j < 10^{-6}$. Fixing $\eta = 1.0$, Fig. 2 shows the trajectory of the gap δ_j for varying p values.

4 Adaptive Proximal Bundle Method

As discussed in Sect. 3, the cutting-plane method (i.e., Algorithm 2) is widely used in the proximal bundle method as a subroutine to repeatedly solve the proximal subproblem (3). Since the proximal bundle method uses a more accurate cutting-plane model f_j rather than a linearization as an approximation of the objective function f , it generalizes the subgradient method and is able to work with weaker regularization,

namely larger stepsize η . This explains why both methods have optimal complexity bounds [41, 42], but the proximal bundle method is always more efficient in practice.

For both the subgradient method and the proximal bundle method to have the optimal performance, one needs to carefully select the stepsize η , namely, being small enough for the subgradient method and within a certain (but relatively large) range for the proximal bundle method. In both cases, we need to know the problem-dependent parameters such as α and L_α , which are unknown or hard to estimate in practice. In this section, we develop the APBM based on an adaptive stepsize strategy for (3) and using Algorithm 2 to solve each subproblem (3). We also discuss variants of adaptive subgradient methods and compare them with APBM. For simplicity, we only present the analysis of Hölder smooth functions satisfying (5), while the hybrid functions satisfying (6) can be similarly analyzed using results from Subsection 3.2.

From practical observations [39], the proximal bundle method works well when the number of inner iterations (i.e., those of Algorithm 2 to solve (3)) stays as a constant much larger than 1 (i.e., that of the subgradient method), say 10. Recall from Theorem 3.1 that inner complexity is $\tilde{\mathcal{O}}\left(\eta L_\alpha^{\frac{2}{\alpha+1}} \left(\frac{1}{\delta}\right)^{\frac{1-\alpha}{\alpha+1}} + 1\right)$. Since we do not know α and L_α , we cannot choose a constant stepsize η so that the number of inner iterations is close to a desired number such as 10. Hence, an adaptive stepsize rule is indeed needed.

By carefully examining Proposition 3.1 and Theorem 3.1, we find that the inner complexity is $\tilde{\mathcal{O}}(\beta^{-1} + 1)$ where β is as in (16). Suppose we want to prescribe the number of inner iterations to be close to β_0^{-1} for some $\beta_0 \in (0, 1]$, if $\beta_0 \leq \beta$, then by Proposition 3.1(a), we have

$$(1 + \beta_0)\delta_j \leq \delta_{j-1}. \quad (21)$$

Hence, it suffices to begin with a relatively large η , check (21) to determine whether the η is small enough (i.e., β is large enough), and adjust η (if necessary) by progressively halving it.

Algorithm 3 below is a formal statement of APBM based on the above intuition.

Algorithm 3 Adaptive Proximal Bundle Method

Require: Let $y_0 \in \mathbb{R}^d$, $\eta_0 > 0$, $\beta_0 \in (0, 1]$, and $\varepsilon > 0$ be given.

for $k = 1, 2, \dots$ **do**

 Call Algorithm 2 with $(y, \eta, \delta) = (y_{k-1}, \eta_{k-1}, \varepsilon/2)$ and output $(y_k, \tilde{y}_k) = (x_J, \tilde{x}_J)$.

if (21) is always true in the execution of Algorithm 2, **then**

 set $\eta_k = \eta_{k-1}$;

else

 set $\eta_k = \eta_{k-1}/2$.

end if

end for

The following lemma provides basic results of Algorithm 2 and is the starting point of the analysis of APBM.

Lemma 4.1 *Assume f is convex and L_α -Hölder smooth. The following statements hold for APBM:*

a) for every $k \geq 1$ and $u \in \mathbb{R}^d$, we have

$$2\eta_{k-1}[f(\tilde{y}_k) - f(u)] \leq \|y_{k-1} - u\|^2 - \|y_k - u\|^2 + \eta_{k-1}\varepsilon; \quad (22)$$

b) for any $k \geq 1$, if

$$\eta_{k-1} \leq \frac{1}{2\beta_0} \left(\frac{\alpha+1}{L_\alpha} \right)^{\frac{2}{\alpha+1}} \left(\frac{\varepsilon}{2} \right)^{\frac{1-\alpha}{\alpha+1}}, \quad (23)$$

then $\eta_k = \eta_{k-1}$;

c) $\{\eta_k\}$ is a non-increasing sequence;

d) for every $k \geq 0$,

$$\eta_k \geq \underline{\eta} := \min \left\{ \frac{1}{4\beta_0} \left(\frac{\alpha+1}{L_\alpha} \right)^{\frac{2}{\alpha+1}} \left(\frac{\varepsilon}{2} \right)^{\frac{1-\alpha}{\alpha+1}}, \eta_0 \right\}. \quad (24)$$

Proof a) It follows from Lemma 3.1(d) that for every $u \in \mathbb{R}^d$

$$2\eta[f(\tilde{x}_J) - f(x)] \leq 2\eta\delta + \|u - y\|^2 - \|x_J - u\|^2.$$

Noting from step 2 of Algorithm 3 that $(\delta, \eta, y, x_J, \tilde{x}_J) = (\varepsilon/2, \eta_{k-1}, y_{k-1}, y_k, \tilde{y}_k)$, which together with the above inequality, implies that (22) holds.

b) It follows from Proposition 3.1(a) and (23) that (21) always holds in the execution of Algorithm 2. In view of step 3 of Algorithm 3, there holds $\eta_k = \eta_{k-1}$.

c) This statement clearly follows from step 3 of Algorithm 3.

d) This statement immediately follows from (b) and step 3 of Algorithm 3. \square

The following theorem gives the total iteration-complexity of APBM.

Theorem 4.1 Assume f is convex and L_α -Hölder smooth. If $\eta_0 \leq \|y_0 - x_*\|^2/\varepsilon$, then the iteration-complexity to obtain an ε -solution to (1) (i.e., a point \hat{x} such that $f(\hat{x}) - \min_{x \in \mathbb{R}^d} f(x) \leq \varepsilon$) is given by

$$\tilde{\mathcal{O}} \left(\frac{L_\alpha^{\frac{2}{\alpha+1}} \|y_0 - x_*\|^2}{\varepsilon^{\frac{2}{\alpha+1}}} + \eta_0 L_\alpha^{\frac{2}{\alpha+1}} \left(\frac{1}{\varepsilon} \right)^{\frac{1-\alpha}{\alpha+1}} \log \left(\frac{\eta_0}{\underline{\eta}} \right) + 1 \right) \quad (25)$$

where $\underline{\eta}$ is as in (24).

Proof Noting from Lemma 4.1 (d) that η_k is bounded from below, we know there exists some $\tilde{\eta} \in [\underline{\eta}, \eta_0]$ such that for some $k_0 \geq 1$, $\eta_k \equiv \tilde{\eta}$ for $k \geq k_0$. Thus, it follows from the assumption that $\eta_0 \leq \|y_0 - x_*\|^2/\varepsilon$ that

$$\underline{\eta} \leq \tilde{\eta} \leq \frac{\|y_0 - x_*\|^2}{\varepsilon}. \quad (26)$$

We consider the worst-case scenario where APBM keeps halving the stepsize until it is stable at $\tilde{\eta}$ and the convergence relies on the conservative stepsize $\tilde{\eta}$. Summing (22) from $k = 1$ to n , we have

$$\begin{aligned} 2 \sum_{k=1}^n \eta_{k-1} \left(\min_{1 \leq k \leq n} f(\tilde{y}_k) - f(u) \right) &\leq 2 \sum_{k=1}^n \eta_{k-1} [f(\tilde{y}_k) - f(u)] \\ &\leq \|y_0 - u\|^2 - \|y_n - u\|^2 + \varepsilon \sum_{k=1}^n \eta_{k-1}. \end{aligned}$$

The above inequality with $u = x_*$, the fact that $\eta_k \leq \eta_0$, and the assumption that $\eta_k \equiv \tilde{\eta}$ for $k \geq k_0$ imply that

$$\|y_n - x_*\|^2 \leq \|y_0 - x_*\|^2 + n\eta_0\varepsilon \quad (27)$$

and

$$\min_{1 \leq k \leq n} f(\tilde{y}_k) - f_* \leq \frac{\|y_0 - x_*\|^2}{2 \sum_{k=1}^n \eta_{k-1}} + \frac{\varepsilon}{2} \leq \frac{\|y_0 - x_*\|^2}{2(n - k_0)\tilde{\eta}} + \frac{\varepsilon}{2}.$$

In order to have $\min_{1 \leq k \leq n} f(\tilde{y}_k) - f_* \leq \varepsilon$, we need

$$n - k_0 = \mathcal{O} \left(\frac{\|y_0 - x_*\|^2}{\tilde{\eta}\varepsilon} + 1 \right). \quad (28)$$

Moreover, it follows from the way η_k is updated in step 3 and Lemma 4.1(d) that

$$k_0 = \mathcal{O} \left(\log \left(\frac{\eta_0}{\tilde{\eta}} \right) + 1 \right) = \mathcal{O} \left(\log \left(\frac{\eta_0}{\underline{\eta}} \right) + 1 \right). \quad (29)$$

Indeed, (27) holds with n replaced by any $k \leq n$ and

$$\|y_k - x_*\|^2 \leq \|y_0 - x_*\|^2 + n\eta_0\varepsilon.$$

It thus follows from (28) and (29) that $\{y_k\}$ is bounded. As a result, using Lemma 3.3, we can derive a uniform bound on δ_1 for every call to Algorithm 2. Now, using Theorem 3.1, we have the iteration-complexity of every call to Algorithm 2 is uniformly bounded by

$$\tilde{\mathcal{O}} \left(\tilde{\eta} L_\alpha^{\frac{2}{\alpha+1}} \left(\frac{1}{\varepsilon} \right)^{\frac{1-\alpha}{\alpha+1}} + 1 \right) \quad (30)$$

for every cycle $k \geq k_0$ and by

$$\tilde{\mathcal{O}} \left(\eta_0 L_\alpha^{\frac{2}{\alpha+1}} \left(\frac{1}{\varepsilon} \right)^{\frac{1-\alpha}{\alpha+1}} + 1 \right) \quad (31)$$

for every cycle $k \leq k_0 - 1$. Hence, multiplying (28) and (30) and using (26) and the definition of $\underline{\eta}$ in (24), we obtain the iteration-complexity

$$\tilde{\mathcal{O}} \left(\frac{L_{\alpha}^{\frac{2}{\alpha+1}} \|y_0 - x_*\|^2}{\varepsilon^{\frac{2}{\alpha+1}}} + 1 \right)$$

for cycles $k \geq k_0$, and multiplying (29) and (31), we obtain the iteration-complexity

$$\tilde{\mathcal{O}} \left(\eta_0 L_{\alpha}^{\frac{2}{\alpha+1}} \left(\frac{1}{\varepsilon} \right)^{\frac{1-\alpha}{\alpha+1}} \log \left(\frac{\eta_0}{\underline{\eta}} \right) + 1 \right)$$

for cycles $k \leq k_0 - 1$. Finally, the total iteration-complexity (25) clearly follows from the above two bounds. \square

We note that the final ε -solution produced by Algorithm 3 is the point \tilde{y}_k that achieves $\min_{1 \leq k \leq n} f(\tilde{y}_k)$. This differs from the last-iterate convergence observed in the smooth case, since the objective function f here is Hölder smooth and may include the nonsmooth case (i.e., $\alpha = 0$).

Discussion on other universal methods Several universal methods based on the backtracking line-search procedure have been studied in the literature. Paper [53] considers the same Hölder smooth problem (with an additional hybrid function h) as in this paper. To finds an ε -solution of (1), the universal primal gradient method proposed in [53] starts from an initial pair (\hat{x}_0, η_0) and in the $(j + 1)$ -th iteration searches for a pair (x_η, η) satisfying a condition

$$f(x_\eta) - \ell_f(x_\eta; \hat{x}_j) - \frac{1}{2\eta} \|x_\eta - \hat{x}_j\|^2 \leq \frac{\varepsilon}{2}, \quad (32)$$

where $\ell_f(u; v) = f(v) + \langle f'(v), u - v \rangle$ and

$$x_\eta = \underset{u \in \mathbb{R}^n}{\operatorname{argmin}} \left\{ \ell_f(u; \hat{x}_j) + h(u) + \frac{1}{2\eta} \|u - \hat{x}_j\|^2 \right\}. \quad (33)$$

If the condition (32) is not satisfied, then the method rejects the pair, sets $\eta \leftarrow \eta/2$, and updates x_η as in (33) with the new η , otherwise, it accepts the pair and sets $(\hat{x}_{j+1}, \eta_{j+1}) = (x_\eta, \eta)$. Two other universal methods are developed in [53], namely, the universal dual gradient method and the universal fast gradient method. Following [53], paper [21] extends the universal fast gradient method to the case of hybrid functions (6). Motivated by the bundle-level method of [36], paper [32] proposes two accelerated variants, i.e., the accelerated bundle-level method and the accelerated prox-level method. The parallel bundle method of [9] is also shown to be universal at the price of running multiple threads.

Paper [42] proposes an adaptive composite subgradient (A-CS) method for solving (1) where f satisfies

$$\|f'(x) - f'(y)\| \leq 2M_f + L_f \|x - y\|, \quad \forall x, y \in \mathbb{R}^d. \quad (34)$$

It is shown in Proposition 2.1 of [42] that any function f that satisfies

$$\|f'(x) - f'(y)\| \leq 2M_\alpha + L_\alpha \|x - y\|^\alpha, \quad \forall x, y \in \mathbb{R}^d,$$

for some $\alpha \in [0, 1]$ also satisfies (34) with

$$M_f(\theta) := M_\alpha + \frac{L_\alpha \theta}{2}, \quad L_f(\theta) := L_\alpha \alpha \left(\frac{1 - \alpha}{\theta} \right)^{\frac{1 - \alpha}{\alpha}}$$

for any $\theta > 0$. Hence, the Hölder smooth functions (5) considered in this paper are included in the class of functions satisfying (34). More interestingly, a careful look at A-CS of [42] and the universal primal gradient method of [53] reveals that the two methods are identical.

The universal primal gradient method is essentially an adaptive subgradient method and the convergence of subgradient methods relies on small enough stepsizes, so it is natural to enforce (32) to make the method adaptive. However, the bundle method converges with any constant stepsize η since it guarantees the condition $\delta_j \leq \varepsilon/2$, which is in the same spirit of (32), by the cutting-plane approach (i.e., Algorithm 2) but not by small η . Therefore, it is not necessary to use a small η in every iteration of each call to Algorithm 2. Instead of frequently reducing η , by the introduction of β_0 , APBM develops a way to regulate the complexity of Algorithm 2 and adjust η only when (21) is not always true in the previous call to Algorithm 2. Another difference between APBM and the universal primal gradient method is that the latter rejects all the pairs (x_η, η) until (32) is satisfied, but APBM always accepts the output of Algorithm 2 even if (21) is not true for every iteration in Algorithm 2. Therefore, APBM potentially employs a larger stepsize η than the universal primal gradient method and is thus a more relaxed adaptive method.

Discussion on optimal universal methods The lower complexity bound for solving (1) is shown in [50] to be

$$\mathcal{O} \left(\left(\frac{L_\alpha \|y_0 - x_*\|^{1+\alpha}}{\varepsilon} \right)^{\frac{2}{1+3\alpha}} \right).$$

The well-known Nesterov's accelerated gradient method has been shown in [51] to match the above complexity bound and hence is an optimal method. The accelerated bundle-level method of [32], the universal fast gradient method of [53], and a follow-up work [21] all establish optimal complexity bounds.

On the other hand, the dominant term of bound (25) is its first term and it is only optimal when $\alpha = 0$, i.e., f is L_0 -Lipschitz continuous. Motivated by [51], it

is possible to develop optimal universal methods based on the accelerated gradient method. This requires accelerated schemes in both PPF and Algorithm 2. Paper [46] proposes an accelerated variant of PPF, which is extended by [4, 15, 23] to obtain optimal p -th order methods with convergence rate $\mathcal{O}(k^{-(3p+1)/2})$ for $p \geq 2$.

We finally note that this paper does not aim to develop the optimal complexity of universal methods; rather, it presents an interesting application of our analysis of Algorithm 2 in the context of universal methods.

5 Proximal Sampling Algorithm

Assuming the RGO in the ASF can be realized, the ASF exhibits remarkable convergence properties. It was shown in [33] that Algorithm 1 converges linearly when f is strongly convex. This convergence result is recently improved in [6] under various weaker assumptions on the target distribution $\pi^X \propto \exp(-f)$. Below we present several convergence results established in [6] that will be used in this paper, under the assumptions that π^X is log-concave, or satisfies the log-Sobolev inequality or Poincaré inequality (PI). Recall that a probability distribution ν satisfies PI with constant $C_{\text{PI}} > 0$ ($1/C_{\text{PI}}$ -PI) if for any smooth bounded function $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$,

$$\mathbb{E}_\nu[(\psi - \mathbb{E}_\nu(\psi))^2] \leq C_{\text{PI}} \mathbb{E}_\nu[\|\nabla \psi\|^2].$$

To this end, for two probability distributions $\rho \ll \nu$, we denote by

$$H_\nu(\rho) := \int \rho \log \frac{\rho}{\nu}, \quad \chi_\nu^2(\rho) := \int \frac{\rho^2}{\nu} - 1$$

the *KL divergence* and the *Chi-squared divergence*, respectively. We denote by W_2 the Wasserstein-2 distance

$$W_2^2(\nu, \rho) := \min_{\gamma \in \Pi(\nu, \rho)} \int \|x - y\|^2 d\gamma(x, y),$$

where $\Pi(\nu, \rho)$ represents the set of all couplings between ν and ρ .

Theorem 5.1 ([6, Theorems 2 & 4]) *We denote by ρ_k^X the law of x_k of Algorithm 1 starting from any initial distribution ρ_0^X . Then, the following statements hold:*

- if $\pi^X \propto \exp(-f)$ is log-concave (i.e., f is convex), then $H_{\pi^X}(\rho_k^X) \leq W_2^2(\rho_0^X, \pi^X)/(k\eta)$;*
- if $\pi^X \propto \exp(-f)$ satisfies λ -PI, then $\chi_{\pi^X}^2(\rho_k^X) \leq \chi_{\pi^X}^2(\rho_0^X)/(1 + \lambda\eta)^{2k}$.*

As discussed earlier, to use ASF in sampling problems, we need to realize the RGO with efficient implementations. In the rest of this section, we develop efficient algorithms for RGO associated with the two scenarios of sampling we are interested in, and then combine them with the ASF to establish a proximal algorithm for sampling. The complexity of the proximal algorithm can be obtained by combining the above

convergence results for ASF and the complexity results we develop for RGO. The rest of the section is organized as follows. In Subsection 5.1 we develop an efficient algorithm for RGO associated with Hölder smooth potentials via rejection sampling. This is combined with ASF to obtain an efficient sampling algorithm from Hölder smooth potentials. In Subsection 5.2, we further extend results to the second setting, i.e., hybrid potentials.

5.1 Sampling from Hölder Smooth Potentials

The bottleneck of using the ASF (Algorithm 1) in sampling tasks with general distributions is the availability of RGO implementations. In this subsection, we address this issue for convex Hölder smooth potentials by developing an efficient algorithm for the corresponding RGO.

Our algorithm of RGO for f is based on rejection sampling. We use a special proposal, namely a Gaussian distribution centered at the δ -solution of (11), which is obtained by invoking Algorithm 2. With this proposal and a sufficiently small $\eta > 0$, the expected number of rejection sampling steps to obtain one effective sample turns out to be bounded from above by a dimension-free constant. To bound the complexity of the rejection sampling, we develop a novel technique to estimate a modified Gaussian integral (see Proposition 5.1).

To this end, let J , \tilde{x}_J , x_J be the outputs of Algorithm 2 and define

$$h_1 := \frac{1}{2\eta} \|\cdot - x_J\|^2 + f_y^\eta(\tilde{x}_J) - \delta, \quad (35a)$$

$$h_2 := \frac{1}{2\eta} \|\cdot - x^*\|^2 + \frac{L_\alpha}{\alpha + 1} \|\cdot - x^*\|^{\alpha+1} + f_y^\eta(x^*). \quad (35b)$$

Note that h_2 is only used for analysis and thus the fact it depends on x^* is not an issue. Algorithm 4 describes the implementation of RGO for f based on Algorithm 2 and rejection sampling.

Algorithm 4 RGO Implementation based on Rejection Sampling

```

1. Let  $y \in \mathbb{R}^d$ ,  $\eta > 0$ , and  $\delta > 0$  be given, and run Algorithm 2 to compute  $x_J$  and  $\tilde{x}_J$ .
2. Generate  $X \sim \exp(-h_1(x))$ .
3. Generate  $U \sim \mathcal{U}[0, 1]$ .
if  $U \leq \exp(-f_y^\eta(X) + h_1(X))$ , then
    accept/return  $X$ ;
else
    reject  $X$  and go to step 2.
end if

```

Lemma 5.1 *Assume f is convex and L_α -Hölder smooth. Let f_y^η be as in (11) and h_1 and h_2 be as in (35). Then, for every $x \in \mathbb{R}^d$, we have*

$$h_1(x) \leq f_y^\eta(x) \leq h_2(x). \quad (36)$$

Proof The first inequality in (36) immediately follows from Lemma 3.1(d) and the definition of h_1 in (35a). By the definition of f_y^η in (11) we get

$$\begin{aligned} f_y^\eta(x) - f_y^\eta(x^*) &= f(x) - f(x^*) + \frac{1}{2\eta} \|x - y\|^2 - \frac{1}{2\eta} \|x^* - y\|^2 \\ &= f(x) - f(x^*) + \frac{1}{2\eta} \|x - x^*\|^2 + \frac{1}{\eta} \langle x - x^*, x^* - y \rangle. \end{aligned} \quad (37)$$

It follows from Lemma 3.1(e) and (7) with $(u, v) = (x, x^*)$ that

$$f(x) - f(x^*) + \frac{1}{\eta} \langle x^* - y, x - x^* \rangle \leq \frac{L_\alpha}{\alpha + 1} \|x - x^*\|^{\alpha+1},$$

which together with (37) implies that

$$f_y^\eta(x) - f_y^\eta(x^*) \leq \frac{L_\alpha}{\alpha + 1} \|x - x^*\|^{\alpha+1} + \frac{1}{2\eta} \|x - x^*\|^2.$$

Using the above inequality and the definition of h_2 in (35b), we conclude that the second inequality in (36) holds. \square

From the expression of h_1 in (35a), it is clear that the proposal distribution $\exp(-h_1(x))$ is a Gaussian centered at x_J . To achieve a tight bound on the expected runs of the rejection sampling, we use a function h_2 which is not quadratic; the standard choice of quadratic function does not give as tight results due to the lack of smoothness. To use this h_2 in the complexity analysis, we need to estimate the integral $\int \exp(-h_2)$, which turns out to be a *highly nontrivial task*. Below we establish a technical result on a modified Gaussian integral, which will be used later to bound the integral $\int \exp(-h_2)$ and hence the complexity of the RGO rejection sampling in Algorithm 4.

Proposition 5.1 *Let $\alpha \in [0, 1]$, $\eta > 0$, $a \geq 0$ and $d \geq 1$. If*

$$2a(\eta d)^{(\alpha+1)/2} \leq 1, \quad (38)$$

then

$$\int_{\mathbb{R}^d} \exp\left(-\frac{1}{2\eta} \|x\|^2 - a\|x\|^{\alpha+1}\right) dx \geq \frac{(2\pi\eta)^{d/2}}{2}. \quad (39)$$

Proof Denote $r = \|x\|$, then

$$dx = r^{d-1} dr dS^{d-1},$$

where dS^{d-1} is the surface area of the $(d-1)$ -dimensional unit sphere.

It follows that

$$\int_{\mathbb{R}^d} \exp\left(-\frac{1}{2\eta} \|x\|^2 - sa\|x\|^{\alpha+1}\right) dx = \int_0^\infty \int \exp\left(-\frac{1}{2\eta} r^2 - ar^{\alpha+1}\right) r^{d-1} dr dS^{d-1}$$

$$= \frac{2\pi^{d/2}}{\Gamma\left(\frac{d}{2}\right)} \int_0^\infty \exp\left(-\frac{1}{2\eta}r^2 - ar^{\alpha+1}\right) r^{d-1} dr. \quad (40)$$

In the above equation, we have used the fact that the total surface area of a $(d-1)$ -dimensional unit sphere is $2\pi^{d/2}/\Gamma\left(\frac{d}{2}\right)$ where $\Gamma(\cdot)$ is the gamma function, i.e.,

$$\Gamma(z) = \int_0^\infty t^{z-1} e^{-t} dt. \quad (41)$$

Defining

$$F_{d,\eta}(a) := \int_0^\infty \exp\left(-\frac{1}{2\eta}r^2 - ar^{\alpha+1}\right) r^d dr, \quad (42)$$

to establish (39), it suffices to bound $F_{d-1,\eta}(a)$ from below.

It follows directly from the definition of $F_{d,\eta}$ in (42) that

$$\frac{dF_{d-1,\eta}(a)}{da} = \int_0^\infty \exp\left(-\frac{1}{2\eta}r^2 - ar^{\alpha+1}\right) (-r^{\alpha+1}) r^{d-1} dr = -F_{d+\alpha,\eta}(a).$$

This implies $F_{d,\eta}$ is monotonically decreasing and thus $F_{d+\alpha,\eta}(a) \leq F_{d+\alpha,\eta}(0)$. As a result,

$$\frac{dF_{d-1,\eta}(a)}{da} \geq -F_{d+\alpha,\eta}(0)$$

and therefore,

$$F_{d-1,\eta}(a) \geq F_{d-1,\eta}(0) - aF_{d+\alpha,\eta}(0). \quad (43)$$

Setting $t = r^2/(2\eta)$, we can write

$$\begin{aligned} F_{d,\eta}(0) &= \int_0^\infty \exp\left(-\frac{1}{2\eta}r^2\right) r^d dr = \int_0^\infty e^{-t} (2\eta t)^{\frac{d-1}{2}} \eta dt \\ &= 2^{\frac{d-1}{2}} \eta^{\frac{d+1}{2}} \int_0^\infty e^{-t} t^{\frac{d-1}{2}} dt. \end{aligned} \quad (44)$$

In view of the definition of the gamma function (41), we obtain

$$F_{d,\eta}(0) = 2^{\frac{d-1}{2}} \eta^{\frac{d+1}{2}} \Gamma\left(\frac{d+1}{2}\right). \quad (45)$$

Applying the Wendel's double inequality (56)

yields

$$\frac{\Gamma\left(\frac{d+\alpha+1}{2}\right)}{\Gamma\left(\frac{d}{2}\right)} \leq \left(\frac{d}{2}\right)^{\frac{\alpha+1}{2}}.$$

Using (43), (45), the above inequality and the assumption (38), we have

$$F_{d-1,\eta}(a) \geq F_{d-1,\eta}(0) - aF_{d+\alpha,\eta}(0)$$

$$\begin{aligned}
&= 2^{\frac{d}{2}-1} \eta^{\frac{d}{2}} \Gamma\left(\frac{d}{2}\right) - a 2^{\frac{d+\alpha-1}{2}} \eta^{\frac{d+\alpha+1}{2}} \Gamma\left(\frac{d+\alpha+1}{2}\right) \\
&= 2^{\frac{d}{2}-1} \eta^{\frac{d}{2}} \Gamma\left(\frac{d}{2}\right) \left(1 - a 2^{\frac{\alpha+1}{2}} \eta^{\frac{\alpha+1}{2}} \frac{\Gamma\left(\frac{d+\alpha+1}{2}\right)}{\Gamma\left(\frac{d}{2}\right)}\right) \\
&\geq 2^{\frac{d}{2}-1} \eta^{\frac{d}{2}} \Gamma\left(\frac{d}{2}\right) \left(1 - a(\eta d)^{\frac{\alpha+1}{2}}\right) \geq \frac{1}{4} (2\eta)^{\frac{d}{2}} \Gamma\left(\frac{d}{2}\right).
\end{aligned}$$

The result (39) then follows from the above inequality and (40). \square

We now proceed to show that the number of rejections in Algorithm 4 is bounded from above by a small constant when δ is properly chosen. In particular, as shown in Proposition 5.2, it only gets worse by a factor of $\exp(\delta)$ and the factor does not depend on the dimension d . Hence, the implementation of RGO for f is computationally efficient in practice.

Proposition 5.2 *Assume f is convex and L_α -Hölder smooth. If*

$$\eta \leq \frac{(\alpha+1)^{\frac{2}{\alpha+1}}}{(2L_\alpha)^{\frac{2}{\alpha+1}} d}, \quad (46)$$

then the expected number of iterations in the rejection sampling of Algorithm 4 is at most $2 \exp(\delta)$.

Proof It is a well-known result for rejection sampling that $X \sim \pi^{X|Y}(x \mid y)$ and the probability that X is accepted is

$$\mathbb{P}\left(U \leq \frac{\exp(-f_y^\eta(X))}{\exp(-h_1(X))}\right) = \frac{\int_{\mathbb{R}^d} \exp(-f_y^\eta(x)) dx}{\int_{\mathbb{R}^d} \exp(-h_1(x)) dx}. \quad (47)$$

If follows directly from the definition of h_2 in (35b) that

$$\int_{\mathbb{R}^d} \exp(-h_2(x)) dx = \exp(-f_y^\eta(x^*)) \int_{\mathbb{R}^d} \exp\left(-\frac{1}{2\eta} \|x - x^*\|^2 - \frac{L_\alpha}{\alpha+1} \|x - x^*\|^{\alpha+1}\right) dx$$

Applying Proposition 5.1 to the above yields

$$\int_{\mathbb{R}^d} \exp(-h_2(x)) dx \geq \exp(-f_y^\eta(x^*)) \frac{(2\pi\eta)^{d/2}}{2}.$$

Note that the condition (38) in Proposition 5.1 holds thanks to (46). By Lemma 5.1, the above inequality leads to

$$\int_{\mathbb{R}^d} \exp(-f_y^\eta(x)) dx \geq \int_{\mathbb{R}^d} \exp(-h_2(x)) dx \geq \exp(-f_y^\eta(x^*)) \frac{(2\pi\eta)^{d/2}}{2}. \quad (48)$$

Using the definition of h_1 in (35a) and Lemma A.1, we have

$$\int_{\mathbb{R}^d} \exp(-h_1(x)) dx = \exp\left(-f_y^\eta(\tilde{x}_J) + \delta\right) (2\pi\eta)^{d/2}. \quad (49)$$

Using (47), (48) and the above identity, we conclude that

$$\mathbb{P}\left(U \leq \frac{\exp(-f_y^\eta(X))}{\exp(-h_1(X))}\right) \geq \frac{1}{2} \exp(-f_y^\eta(x^*) + f_y^\eta(\tilde{x}_J) - \delta) \geq \frac{1}{2} \exp(-\delta),$$

and the expected number of the iterations is

$$\frac{1}{\mathbb{P}\left(U \leq \frac{\exp(-f_y^\eta(X))}{\exp(-h_1(X))}\right)} \leq 2 \exp(\delta).$$

□

We finally bound the total complexity to sample from a log-concave distribution ν in (2) with a Hölder smooth potential f . We combine our efficient algorithm (Algorithm 4) of RGO for Hölder smooth potentials and the convergent results for ASF, namely Theorem 5.1, to achieve this goal.

Theorem 5.2 *Assume f is convex and L_α -Hölder smooth, then Algorithm 1, initialized with ρ_0^X and stepsize $\eta \asymp 1/(L_\alpha^{\frac{2}{\alpha+1}} d)$, using Algorithm 4 as an RGO has the iteration-complexity bound*

$$\mathcal{O}\left(\frac{L_\alpha^{\frac{2}{\alpha+1}} d W_2^2(\rho_0^X, \nu)}{\varepsilon}\right)$$

to achieve ε error to the target $\nu \propto \exp(-f)$ in terms of KL divergence. Each RGO requires $\tilde{\mathcal{O}}\left(\frac{1}{d} \left(\frac{1}{\delta}\right)^{\frac{1-\alpha}{\alpha+1}} + 1\right)$ subgradient evaluations of f and $2 \exp(\delta)$ rejection steps in expectation. Moreover, if ν satisfies PI with constant $C_{\text{PI}} > 0$,

then the iteration-complexity bound to achieve ε error in terms of Chi-squared divergence is

$$\tilde{\mathcal{O}}\left(C_{\text{PI}} L_\alpha^{\frac{2}{\alpha+1}} d\right).$$

Proof The results follow directly from Theorem 5.1, Theorem 3.1 and Proposition 5.2 with the choice of stepsize $\eta \asymp 1/(L_\alpha^{\frac{2}{\alpha+1}} d)$. □

5.2 Sampling from Hybrid Potentials

In this subsection, we consider sampling from a log-concave distribution $\nu \propto \exp(-f(x))$ associated with a hybrid potential f satisfying (6). This setting is a generalization of the Hölder smooth setting studied in the previous sections. It turns out

that both Algorithm 1 and the implementation for RGO, Algorithm 4, developed for Hölder smooth sampling can be applied directly to this general setting with properly chosen stepsizes. Below, we extend the analysis in Subsection 5.1 to the hybrid setting and establish corresponding complexity results.

The following lemma is a counterpart of Lemma 5.1 in the hybrid setting. Its proof is given in Appendix B.

Lemma 5.2 *Assume f is convex and satisfies (6). Define*

$$h_2(x) := \frac{1}{2\eta} \|x - x^*\|^2 + \sum_{i=1}^n \frac{L_{\alpha_i}}{\alpha_i + 1} \|x - x^*\|^{\alpha_i + 1} + f_y^\eta(x^*). \quad (50)$$

Then, $h_2(x) \geq f_y^\eta(x)$ for every $x \in \mathbb{R}^d$.

The next result is an analogue of the modified Gaussian integral in Proposition 5.1.

Proposition 5.3 *Let $\alpha_i \in [0, 1]$, $a_i \geq 0$, $\eta > 0$, and $d \geq 1$. If*

$$\eta d \sum_{i=1}^n a_i^{\frac{2}{\alpha_i + 1}} \leq 1, \quad (51)$$

then

$$\int_{\mathbb{R}^d} \exp \left(-\frac{1}{2\eta} \|x\|^2 - \sum_{i=1}^n a_i \|x\|^{\alpha_i + 1} \right) dx \geq (2\pi\eta)^{\frac{d}{2}} \exp \left(-\frac{1}{2} + \frac{\sum_{i=1}^n (\alpha_i - 1)}{4} \right). \quad (52)$$

Proof Using the Young's inequality $st \leq s^p/p + t^q/q$ with

$$s = a 2^{\frac{1-\alpha}{2}} \|x\|^{\alpha+1}, \quad t = \frac{1}{2^{\frac{1-\alpha}{2}}}, \quad p = \frac{2}{\alpha+1}, \quad q = \frac{2}{1-\alpha},$$

we obtain

$$a \|x\|^{\alpha+1} \leq (\alpha+1) a^{\frac{2}{\alpha+1}} 2^{\frac{-2\alpha}{\alpha+1}} \|x\|^2 + \frac{1-\alpha}{4} \leq a^{\frac{2}{\alpha+1}} \|x\|^2 + \frac{1-\alpha}{4},$$

where the second inequality is due to the fact that $(\alpha+1) 2^{\frac{-2\alpha}{\alpha+1}} \leq 1$ for $\alpha \in [0, 1]$. Hence, the above inequality generalizes to

$$\sum_{i=1}^n a_i \|x\|^{\alpha_i + 1} \leq \sum_{i=1}^n a_i^{\frac{2}{\alpha_i + 1}} \|x\|^2 + \sum_{i=1}^n \frac{1-\alpha_i}{4}.$$

This inequality and Lemma A.1 imply that

$$\int_{\mathbb{R}^d} \exp \left(-\frac{1}{2\eta} \|x\|^2 - \sum_{i=1}^n a_i \|x\|^{\alpha_i + 1} \right) dx$$

$$\begin{aligned}
&\geq \int_{\mathbb{R}^d} \exp \left(-\frac{1}{2\eta} \|x\|^2 - \sum_{i=1}^n a_i^{\frac{2}{\alpha_i+1}} \|x\|^2 - \sum_{i=1}^n \frac{1-\alpha_i}{4} \right) dx \\
&= \exp \left(\frac{\sum_{i=1}^n (\alpha_i - 1)}{4} \right) \int_{\mathbb{R}^d} \exp \left(-\frac{1}{2\tilde{\eta}} \|x\|^2 \right) dx \\
&= \exp \left(\frac{\sum_{i=1}^n (\alpha_i - 1)}{4} \right) (2\pi\tilde{\eta})^{\frac{d}{2}}
\end{aligned} \tag{53}$$

where

$$\frac{1}{\tilde{\eta}} = \frac{1}{\eta} + \sum_{i=1}^n a_i^{\frac{2}{\alpha_i+1}}. \tag{54}$$

It follows from (51) that

$\tilde{\eta} \geq (1 + \frac{1}{d})^{-1} \eta$. Plugging this inequality into (53), we have

$$\begin{aligned}
\int_{\mathbb{R}^d} \exp \left(-\frac{1}{2\eta} \|x\|^2 - a\|x\|^{\alpha+1} \right) dx \\
\geq (2\pi\eta)^{\frac{d}{2}} \left(1 + \frac{1}{d} \right)^{-\frac{d}{2}} \exp \left(\frac{\sum_{i=1}^n (\alpha_i - 1)}{4} \right) \\
\geq (2\pi\eta)^{\frac{d}{2}} \exp \left(-\frac{1}{2} + \frac{\sum_{i=1}^n (\alpha_i - 1)}{4} \right),
\end{aligned}$$

where in the second inequality, we use the fact that

$$\left(1 + \frac{1}{d} \right)^{\frac{d}{2}} \leq \exp \left(\frac{1}{2} \right).$$

□

With Lemma 5.2 and Proposition 5.3 in hand, we can bound the complexity of Algorithm 4 as follows. The proof is postponed to Appendix B.

Proposition 5.4 *If stepsize η satisfies*

$$\eta d \sum_{i=1}^n \left(\frac{L_{\alpha_i}}{\alpha_i + 1} \right)^{\frac{2}{\alpha_i+1}} \leq 1, \tag{55}$$

then rejection steps in Algorithm 4 take at most $\exp \left(\delta + \frac{1}{2} + \frac{\sum_{i=1}^n (1-\alpha_i)}{4} \right)$ iterations in expectation.

Through the above arguments, we show that Algorithm 4 designed for Hölder smooth potentials is equally effective for hybrid potentials satisfying (6). Combining Proposition 5.4 and Theorem 3.2 with the convergence results for ASF, we obtain the following iteration-complexity bounds for sampling from hybrid potentials. The proof is similar to that of Theorem 5.2 and is thus omitted.

Theorem 5.3 Assume f is a convex and satisfies (6). Consider Algorithm 1, initialized with ρ_0^X and stepsize η satisfies (55), using Algorithm 4 as a RGO. Each RGO requires $\tilde{\mathcal{O}}\left(\frac{1}{d\delta} + 1\right)$ subgradient evaluations of f and in expectation $\exp\left(\delta + \frac{1}{2} + \frac{\sum_{i=1}^n(1-\alpha_i)}{4}\right)$ rejection steps. The total complexity of Algorithm 1 to achieve ε error in terms of KL divergence is

$$\mathcal{O}\left(\frac{\sum_{i=1}^n \left(\frac{L_{\alpha_i}}{\alpha_i+1}\right)^{\frac{2}{\alpha_i+1}} dW_2^2(\rho_0^X, v)}{\varepsilon}\right).$$

Moreover, if v satisfies PI with constant $C_{\text{PI}} > 0$, then total complexity to achieve ε error in terms of Chi-squared divergence is

$$\tilde{\mathcal{O}}\left(\sum_{i=1}^n \left(\frac{L_{\alpha_i}}{\alpha_i+1}\right)^{\frac{2}{\alpha_i+1}} dC_{\text{PI}}\right).$$

6 Conclusions

In this paper, we study proximal algorithms for both optimization and sampling lacking smoothness. We first establish the complexity bounds of the regularized cutting-plane method for solving proximal subproblem (3), where f is convex and satisfies either (5) (Hölder smooth) or (6) (hybrid). This efficient implementation gives an approximate solution to the proximal map in optimization, which is the core of both proximal optimization and sampling algorithms.

For optimization, we develop APBM using a novel adaptive stepsize strategy in the proximal point method and the approximate proximal map to solve each proximal subproblem. The proposed APBM is a universal method as it does not require any problem-dependent parameters as input.

For sampling, we propose an efficient method based on rejection sampling and the approximate proximal map to realize the RGO, which is a proximal sampling oracle. Finally, combining the sampling complexity of RGO and the complexity bounds of ASF, which is a counterpart of the proximal point method in sampling, we establish the complexity bounds of the proximal sampling algorithm in both Hölder smooth and hybrid settings.

This paper provides a unified perspective to study proximal optimization and sampling algorithms, while many other interesting questions remain open. First, APBM is only optimal when $\alpha = 0$, i.e., f is Lipschitz continuous. We are interested in developing a universal method that is optimal for any $\alpha \in [0, 1]$. One possible direction is to incorporate the acceleration technique into both the regularized cutting-plane method and the PPF. Second, as acceleration methods are widely used in optimization to obtain optimal performance, accelerated proximal sampling algorithms are less explored. It is worth investigating a counterpart of the accelerated proximal point

method [46] in sampling. Finally, we develop APBM as a universal method for non-smooth optimization, and it would be equally important to design a universal method for sampling.

A Technical Results

This section collects technical results that are useful in the paper.

Lemma A.1 (*Gaussian integral*) *For any $\eta > 0$,*

$$\int_{\mathbb{R}^d} \exp\left(-\frac{1}{2\eta}\|x\|^2\right) dx = (2\pi\eta)^{d/2}.$$

The following lemma provides both lower and upper bounds on the ratio of gamma functions. Its proof can be found in [65].

Lemma A.2 (*Wendel's double inequality*) *For $0 < s < 1$ and $t > 0$, the gamma function defined as in (41) satisfies*

$$\left(\frac{t}{t+s}\right)^{1-s} \leq \frac{\Gamma(t+s)}{t^s \Gamma(t)} \leq 1,$$

or equivalently,

$$t^{1-s} \leq \frac{\Gamma(t+1)}{\Gamma(t+s)} \leq (t+s)^{1-s}. \quad (56)$$

Lemma A.3 *Assume f is convex and L_α -Hölder smooth (i.e., satisfying (5), then (7) holds for every $u, v \in \mathbb{R}^d$. Assume f is convex and satisfies (6), then (8) holds for every $u, v \in \mathbb{R}^d$.*

Proof We first consider the case when f is convex and L_α -Hölder smooth. It is easy to see that

$$\begin{aligned} f(u) &= f(v) + \int_0^1 \langle f'(v + \tau(v - u)), u - v \rangle d\tau \\ &= f(v) + \langle f'(v), u - v \rangle + \int_0^1 \langle f'(v + \tau(v - u)) - f'(v), u - v \rangle d\tau. \end{aligned}$$

Using the above identity, the Cauchy-Schwarz inequality, and (5), we have

$$\begin{aligned} f(u) - f(v) - \langle f'(v), u - v \rangle &= \int_0^1 \langle f'(v + \tau(v - u)) - f'(v), u - v \rangle d\tau \\ &\leq \int_0^1 \|f'(v + \tau(v - u)) - f'(v)\| \|u - v\| d\tau \\ &\leq \int_0^1 L_\alpha \tau^\alpha \|u - v\|^{\alpha+1} d\tau = \frac{L_\alpha}{\alpha+1} \|u - v\|^{\alpha+1}. \end{aligned}$$

Hence, (7) holds. More generally, if f satisfies (6), then (8) follows the same argument. \square

Lemma A.4 Consider $\phi(t) = |t|^p$ and $\phi'(t) = p \operatorname{sign}(t) |t|^{p-1}$ for $t \in \mathbb{R}$ and some $p \in [1, 2]$. Then, for any $u, v \in \mathbb{R}$, we have

$$|\phi'(u) - \phi'(v)| \leq p 2^{2-p} |u - v|^{p-1}.$$

Proof Let $r := p - 1 \in [0, 1]$. We consider the following two cases and prove

$$|\phi'(u) - \phi'(v)| \leq p 2^{1-r} |u - v|^r.$$

Case 1: $uv \geq 0$. Here, $\operatorname{sign}(u) = \operatorname{sign}(v)$, so

$$|\phi'(u) - \phi'(v)| = p | |u|^r - |v|^r |.$$

Without loss of generality, assume $a = |u| \geq b = |v| \geq 0$. By the subadditivity of the function $x \mapsto x^r$ with $r \in [0, 1]$, we have $(a - b)^r \geq a^r - b^r$ for all $a \geq b \geq 0$. Hence

$$| |u|^r - |v|^r | = a^r - b^r \leq (a - b)^r = | |u| - |v| |^r = |u - v|^r.$$

Therefore,

$$|\phi'(u) - \phi'(v)| \leq p |u - v|^r.$$

Case 2: $uv < 0$. In this case, $\operatorname{sign}(u) = -\operatorname{sign}(v)$, so

$$|\phi'(u) - \phi'(v)| = p (|u|^r + |v|^r).$$

By the concavity of $x \mapsto x^r$, we have

$$|u|^r + |v|^r \leq 2^{1-r} (|u| + |v|)^r = 2^{1-r} |u - v|^r,$$

and thus

$$|\phi'(u) - \phi'(v)| \leq p 2^{1-r} |u - v|^r.$$

The conclusion immediately follows from the above two cases. \square

B Missing Proofs in Subsection 5.2

Proof of Lemma 5.2 It follows from the same argument as in the proof of Lemma 5.1 that (37) holds. Using (37), Lemma 3.1(e), and (8) with $(u, v) = (x, x^*)$, we conclude that

$$f_y^\eta(x) - f_y^\eta(x^*) \leq \sum_{i=1}^n \frac{L_{\alpha_i}}{\alpha_i + 1} \|x - x^*\|^{\alpha_i + 1} + \frac{1}{2\eta} \|x - x^*\|^2.$$

The lemma immediately follows from the above inequality and the definition of h_2 in (50). \square

Proof of Proposition 5.4 It follows directly from the definition of h_2 in (50) that

$$\begin{aligned} \int_{\mathbb{R}^d} \exp(-h_2(x)) dx &= \exp(-f_y^\eta(x^*)) \int_{\mathbb{R}^d} \exp\left(-\frac{1}{2\eta}\|x - x^*\|^2\right. \\ &\quad \left.- \sum_{i=1}^n \frac{L_{\alpha_i}}{\alpha_i + 1} \|x - x^*\|^{\alpha_i + 1}\right) dx. \end{aligned}$$

It is easy to see that (55) implies that (51) holds with $a_i = \frac{L_{\alpha_i}}{\alpha_i + 1}$. Hence, by Proposition 5.3, we have (5.3) holds with $a_i = \frac{L_{\alpha_i}}{\alpha_i + 1}$, i.e.,

$$\begin{aligned} &\int_{\mathbb{R}^d} \exp\left(-\frac{1}{2\eta}\|x - x^*\|^2 - \sum_{i=1}^n \frac{L_{\alpha_i}}{\alpha_i + 1} \|x - x^*\|^{\alpha_i + 1}\right) dx \\ &\geq (2\pi\eta)^{\frac{d}{2}} \exp\left(-\frac{1}{2} + \frac{\sum_{i=1}^n (\alpha_i - 1)}{4}\right). \end{aligned}$$

The above two inequalities and Lemma 5.2 imply that

$$\begin{aligned} \int_{\mathbb{R}^d} \exp(-f_y^\eta(x)) dx &\geq \int_{\mathbb{R}^d} \exp(-h_2(x)) dx \geq (2\pi\eta)^{\frac{d}{2}} \\ &\exp\left(-f_y^\eta(x^*) - \frac{1}{2} + \frac{\sum_{i=1}^n (\alpha_i - 1)}{4}\right). \end{aligned}$$

As in the proof of Proposition 5.2, (47) and (49) hold. Using (47), (49), and the above inequality, we have

$$\mathbb{P}\left(U \leq \frac{\exp(-f_y^\eta(X))}{\exp(-h_1(X))}\right) \geq \exp\left(f_y^\eta(\tilde{x}_J) - f_y^\eta(x^*) - \delta - \frac{1}{2} + \frac{\sum_{i=1}^n (\alpha_i - 1)}{4}\right).$$

The above inequality and the fact that $f_y^\eta(\tilde{x}_J) \geq f_y^\eta(x^*)$ immediately imply that

$$\frac{1}{\mathbb{P}\left(U \leq \frac{\exp(-f_y^\eta(X))}{\exp(-h_1(X))}\right)} \leq \exp\left(\delta + \frac{1}{2} + \frac{\sum_{i=1}^n (1 - \alpha_i)}{4}\right).$$

□

Acknowledgements The first author was partially supported by GIDS-AI seed funding and AFOSR grant FA9550-25-1-0182. The second author was supported by NSF under grants 1942523 and 2008513.

Data Availability There is no data related to this publication.

References

1. Beck, A.: First-Order Methods in Optimization, vol. 25. SIAM, USA (2017)
2. Bertsimas, D., Vempala, S.: Solving convex programs by random walks. *J. ACM (JACM)* **51**(4), 540–556 (2004)
3. Bou-Rabee, N., Hairer, M.: Nonasymptotic mixing of the MALA algorithm. *IMA J. Numer. Anal.* **33**(1), 80–110 (2013)
4. Bubeck, S., Jiang, Q., Lee, Y.T., Li, Y., Sidford, A.: Near-optimal method for highly smooth convex optimization. In: Alina Beygelzimer., Daniel Hsu. (eds.) *Proceedings of the Thirty-Second Conference on Learning Theory*, vol. 99 of *Proceedings of Machine Learning Research*, pp. 492–507. PMLR, (2019)
5. Chambolle, A., Pock, T.: A first-order primal-dual algorithm for convex problems with applications to imaging. *J. Math. Imag. Vision* **40**(1), 120–145 (2011)
6. Chen, Y., Chewi, S., Salim, A., Wibisono, A.: Improved analysis for a proximal algorithm for sampling. In: Po-Ling Loh., Maxim Raginsky. (eds.) *Proceedings of Thirty Fifth Conference on Learning Theory of Proceedings of Machine Learning Research*, vol. 178, pp. 2984–3014. PMLR, (2022)
7. Dalalyan, A.S.: Theoretical guarantees for approximate sampling from smooth and log-concave densities. *J. Royal Statistical Soc.: Ser. B (Statistical Methodol.)* **79**(3), 651–676 (2017)
8. de Oliveira, W., Sagastizábal, C., Lemaréchal, C.: Convex proximal bundle methods in depth: a unified analysis for inexact oracles. *Math. Program.* **148**(1–2), 241–277 (2014)
9. Díaz, M., Grimmer, B.: Optimal convergence rates for the proximal bundle method. *SIAM J. Optim.* **33**(2), 424–454 (2023)
10. Durmus, A., Majewski, S., Miasojedow, B.: Analysis of Langevin Monte Carlo via convex optimization. *J. Mach. Learn. Res.* **20**(1), 2666–2711 (2019)
11. Durmus, A., Moulines, E., Pereyra, M.: Efficient Bayesian computation by proximal Markov Chain Monte Carlo: when Langevin meets Moreau. *SIAM J. Imag. Sci.* **11**(1), 473–506 (2018)
12. Dyer, M., Frieze, A., Kannan, R.: A random polynomial-time algorithm for approximating the volume of convex bodies. *J. ACM (JACM)* **38**(1), 1–17 (1991)
13. Fan, J., Yuan, B., Chen, Y.: Improved dimension dependence of a proximal algorithm for sampling. In: Gergely Neu., Lorenzo Rosasco. (eds.) *Proceedings of Thirty Sixth Conference on Learning Theory of Proceedings of Machine Learning Research*, vol. 195, pp. 1473–1521. PMLR, (2023)
14. Frangioni, A.: Generalized bundle methods. *SIAM J. Optim.* **13**(1), 117–156 (2002)
15. Gasnikov, A., Dvurechensky, P., Gorbunov, E., Vorontsova, E., Selikhanych, D., Uribe, C.A.: Optimal tensor methods in smooth convex and uniformly convex optimization. In: Alina Beygelzimer., Daniel Hsu. (eds.) *Proceedings of the Thirty-Second Conference on Learning Theory of Proceedings of Machine Learning Research*, vol. 99, pp. 1374–1391. PMLR, (2019)
16. Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B.: *Bayesian Data Analysis*. CRC Press, USA (2013)
17. Geman, S., Geman, D.: Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell.* **6**, 721–741 (1984)
18. Gopi, S., Lee, Y.T., Liu, D.: Private convex optimization via exponential mechanism. In: Po-Ling Loh., Maxim Raginsky. (eds.) *Proceedings of Thirty Fifth Conference on Learning Theory of Proceedings of Machine Learning Research*, vol. 178, pp. 1948–1989. PMLR, (2022)
19. Goulart, P.J., Chen, Y.: Clarabel: An interior-point solver for conic programs with quadratic objectives. Available on [arXiv:2405.12762](https://arxiv.org/abs/2405.12762), (2024)
20. Grenander, U., Miller, M.I.: Representations of knowledge in complex systems. *J. Roy. Stat. Soc.: Ser. B (Methodol.)* **56**(4), 549–581 (1994)
21. Grimmer, B.: On optimal universal first-order methods for minimizing heterogeneous sums. *Optim. Lett.* **18**(2), 427–445 (2024)
22. Innes, M.: Don't unroll adjoint: Differentiating ssa-form programs. Available on [arXiv:1810.07951](https://arxiv.org/abs/1810.07951), (2018)
23. Jiang, B., Wang, H., Zhang, S.: An optimal high-order tensor method for convex optimization. *Math. Oper. Res.* **46**(4), 1390–1412 (2021)
24. Jordan, R., Kinderlehrer, D., Otto, F.: The variational formulation of the fokker-planck equation. *SIAM J. Math. Anal.* **29**(1), 1–17 (1998)
25. Kalai, A.T., Vempala, S.: Simulated annealing for convex optimization. *Math. Oper. Res.* **31**(2), 253–266 (2006)
26. Kannan, R., Lovász, L., Simonovits, M.: Random walks and an $O^*(n^5)$ volume algorithm for convex bodies. *Random Struct. Algorithms* **11**(1), 1–50 (1997)

27. Kiwiel, K.C.: Efficiency of proximal bundle methods. *J. Optim. Theory Appl.* **104**(3), 589–603 (2000)

28. Kong, W., Melo, J.G., Monteiro, R.D.: Complexity of a quadratic penalty accelerated inexact proximal point method for solving linearly constrained nonconvex composite programs. *SIAM J. Optim.* **29**(4), 2566–2593 (2019)

29. Kong, W., Monteiro, R.D.C.: An accelerated inexact proximal point method for solving nonconvex-concave min-max problems. *SIAM J. Optim.* **31**(4), 2558–2585 (2021)

30. Korpelevič, G.M.: An extragradient method for finding saddle points and for other problems. *Èkonom. i Mat. Metody* **12**(4), 747–756 (1976)

31. Krauth, W.: *Statistical Mechanics: Algorithms and Computations*, vol. 13. OUP Oxford, Oxford (2006)

32. Lan, G.: Bundle-level type methods uniformly optimal for smooth and nonsmooth convex optimization. *Math. Program.* **149**(1–2), 1–45 (2015)

33. Lee, Y.T., Shen, R., Tian, K.: Structured logconcave sampling with a restricted gaussian oracle. In: Mikhail Belkin, Samory Kpotufe, (eds.) *Proceedings of Thirty Fourth Conference on Learning Theory of Proceedings of Machine Learning Research*, vol. 134, pp. 2993–3050. PMLR, (2021)

34. Lemaréchal, C.: An extension of davidon methods to non differentiable problems. In: Balinski, M.L., Wolfe, P. (eds.) *Nondifferentiable Optimization*, pp. 95–109. Springer, Berlin (1975)

35. Lemaréchal, C.: Nonsmooth optimization and descent methods. Technical Report, (1978)

36. Lemaréchal, C., Nemirovski, A., Nesterov, Y.: New variants of bundle methods. *Math. Program.* **69**(1–3), 111–147 (1995)

37. Liang, J., Chen, Y.: A proximal algorithm for sampling from non-smooth potentials. In: 2022 Winter Simulation Conference (WSC), pp. 3229–3240, IEEE, (2022)

38. Liang, J., Chen, Y.: A proximal algorithm for sampling. *Transactions on Machine Learning Research*, (2023)

39. Liang, J., Guigues, V., Monteiro, R.D.C.: A single cut proximal bundle method for stochastic convex composite optimization. *Math. Program.* **208**(1), 173–208 (2024)

40. Liang, J., Monteiro, R.D.C.: A doubly accelerated inexact proximal point method for nonconvex composite optimization problems. Available on [arXiv:1811.11378](https://arxiv.org/abs/1811.11378), (2018)

41. Liang, J., Monteiro, R.D.C.: A proximal bundle variant with optimal iteration-complexity for a large range of prox stepsizes. *SIAM J. Optim.* **31**(4), 2955–2986 (2021)

42. Liang, J., Monteiro, R.D.C.: A unified analysis of a class of proximal bundle methods for solving hybrid convex composite optimization problems. *Math. Oper. Res.* **49**(2), 832–855 (2024)

43. Liang, J., Monteiro, R.D.C., Zhang, H.: Proximal bundle methods for hybrid weakly convex composite optimization problems. Available on [arXiv:2303.14896](https://arxiv.org/abs/2303.14896), (2023)

44. Martinet, B.: Regularisation d'inequations variationnelles par approximations successives. *Revue Française d'Informatique et de Recherche Opérationnelle* **4**, 154–159 (1970)

45. Mifflin, R.: A modification and an extension of Lemaréchal's algorithm for nonsmooth minimization. In: *Nondifferential and Variational Techniques in Optimization*, pp. 77–90. Springer, Berlin, (1982)

46. Monteiro, R.D.C., Svaiter, B.F.: An accelerated hybrid proximal extragradient method for convex optimization and its implications to second-order methods. *SIAM J. Optim.* **23**(2), 1092–1125 (2013)

47. Monteiro, R.D.C., Svaiter, B.F.: Iteration-complexity of block-decomposition algorithms and the alternating direction method of multipliers. *SIAM J. Optim.* **23**(1), 475–507 (2013)

48. Mou, W., Flammarion, N., Wainwright, M.J., Bartlett, P.L.: An efficient sampling algorithm for non-smooth composite potentials. *J. Mach. Learn. Res.* **23**(233), 1–50 (2022)

49. Neal, R.M.: MCMC using Hamiltonian dynamics. *Handbook Markov Chain Monte Carlo* **2**(11), 2 (2011)

50. Nemirovski, A., Yudin, D.B.: *Problem Complexity and Method Efficiency in Optimization*. Wiley, USA (1983)

51. Nemirovskii, A.S., Nesterov, Y.E.: Optimal methods of smooth convex minimization. *USSR Comput. Math. Math. Phys.* **25**(2), 21–30 (1985)

52. Nesterov, Yu.: Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM J. Optim.* **22**(2), 341–362 (2012)

53. Nesterov, Yu.: Universal gradient methods for convex optimization problems. *Math. Program.* **152**(1–2), 381–404 (2015)

54. Parikh, N., Boyd, S.: Proximal algorithms. *Found. Trends Optim.* **1**(3), 127–239 (2014)

55. Parisi, G.: Correlation functions and computer simulations. *Nucl. Phys. B* **180**(3), 378–384 (1981)

56. Roberts, G.O., Stramer, O.: Langevin diffusions and Metropolis-Hastings algorithms. *Methodol. Comput. Appl. Probab.* **4**(4), 337–357 (2002)

57. Roberts, G.O., Tweedie, R.L.: Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli* **2**(4), 341–363 (1996)
58. Rockafellar, R.T.: Augmented Lagrangians and applications of the proximal point algorithm in convex programming. *Math. Oper. Res.* **1**(2), 97–116 (1976)
59. Rockafellar, R.T.: Monotone operators and the proximal point algorithm. *SIAM J. Control. Optim.* **14**(5), 877–898 (1976)
60. Ruszczyński, A.: Nonlinear Optimization. Princeton University Press, NJ (2011)
61. Salim, A., Richtarik, P.: Primal dual interpretation of the proximal stochastic gradient langevin algorithm. In: Larochelle, H., Ranzato, M.A., Hadsell, R., Balcan, M.-F., Lin, H.-T. (eds.) *Advances in Neural Information Processing Systems*, vol. 33, pp. 3786–3796. Curran Associates Inc., USA (2020)
62. Shen, R., Tian, K., Lee, Y.T.: Composite logconcave sampling with a restricted Gaussian oracle. Available on [arXiv:2006.05976](https://arxiv.org/abs/2006.05976), (2020)
63. Sites, J.W., Marshall, J.C.: Delimiting species: a renaissance issue in systematic biology. *Trends Ecol. Evol.* **18**(9), 462–470 (2003)
64. van Ackooij, W., Berge, V.: Wellington de Oliveira, and Claudia Sagastizábal. Probabilistic optimization via approximate p-efficient points and bundle methods. *Comput. Oper. Res.* **77**, 177–193 (2017)
65. Wendel, J.G.: Note on the gamma function. *Am. Math. Mon.* **55**(9), 563–564 (1948)
66. Wibisono, A.: Sampling as optimization in the space of measures: The langevin dynamics as a composite optimization problem. In: Sebastian Bubeck, Vianney Perchet, Philippe Rigollet. (eds.) *Proceedings of the 31st Conference On Learning Theory of Proceedings of Machine Learning Research*, vol.75, pp. 2093–3027, PMLR, (2018)
67. Wolfe, P.: A method of conjugate subgradients for minimizing nondifferentiable functions. In: Balinski, M.L., Wolfe, P. (eds.) *Nondifferentiable Optimization*, pp. 145–173. Springer, Berlin (1975)
68. Yang, Z., Zhang, Y., Chen, Y., Wang, Z.: Variational transport: A convergent particle-based algorithm for distributional optimization. Available on [arXiv:2012.11554](https://arxiv.org/abs/2012.11554), (2020)
69. Yu, D., Ruszczyński, A.: Rate of convergence of the bundle method. *J. Optim. Theory Appl.* **173**(3), 908–922 (2017)
70. Yuan, B., Fan, J., Liang, J., Wibisono, A., Chen, Y.: On a class of gibbs sampling over networks. In: Gergely Neu and Lorenzo Rosasco. (eds.) *Proceedings of Thirty Sixth Conference on Learning Theory of Proceedings of Machine Learning Research*, vol.195, pp.5754–5780, PMLR, (2023)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.