
Stochastic Proximal Langevin Algorithm: Potential Splitting and Nonasymptotic Rates

Adil Salim Dmitry Kovalev Peter Richtárik*
King Abdullah University of Science and Technology, Thuwal, Saudi Arabia

Abstract

We propose a new algorithm—Stochastic Proximal Langevin Algorithm (SPLA)—for sampling from a log concave distribution. Our method is a generalization of the Langevin algorithm to potentials expressed as the sum of one stochastic smooth term and multiple stochastic nonsmooth terms. In each iteration, our splitting technique only requires access to a stochastic gradient of the smooth term and a stochastic proximal operator for each of the nonsmooth terms. We establish nonasymptotic sublinear and linear convergence rates under convexity and strong convexity of the smooth term, respectively, expressed in terms of the KL divergence and Wasserstein distance. We illustrate the efficiency of our sampling technique through numerical simulations on a Bayesian learning task.

1 Introduction

Many applications in the field of Bayesian machine learning require to sample from a probability distribution μ^* with density $\mu^*(x)$, $x \in \mathbb{R}^d$. Due to their scalability, Monte Carlo Markov Chain (MCMC) methods such as Langevin Monte Carlo [48] or Hamiltonian Monte Carlo [28] are popular algorithms to solve such problems. Monte Carlo methods typically generate a sequence of random variables $(x^k)_{k \geq 0}$ with the property that the distribution of x^k approaches μ^* as k grows.

While the theory of MCMC algorithms has remained mainly asymptotic, in recent years the exploration of non-asymptotic properties of such algorithms has led to a renaissance in the field [14, 26, 39, 15, 16, 19, 22, 12, 53, 10, 52]. In particular, if $\mu^*(x) \propto \exp(-U(x))$, where U is a smooth convex function, [14, 19] provide explicit convergence rates for the *Langevin algorithm (LA)*

$$x^{k+1} = x^k - \gamma \nabla U(x^k) + \sqrt{2\gamma} W^k,$$

where $\gamma > 0$ and $(W^k)_{k \geq 0}$ is a sequence of i.i.d. standard Gaussian random variables. The function U , also called the *potential*, enters the algorithm through its gradient.

In optimization, the problem $\min U$ where U is composite, *i.e.* is a sum of nonsmooth terms which must be handled separately, has many instances, see [17, Section 2]. These optimization problems can be seen as a Maximum A Posteriori (MAP) computation of some Bayesian model. Sampling a posteriori in these models allows for a better Bayesian inference [20]. In these cases, the task of sampling a posteriori takes the form of sampling from the target distribution μ^* , where U has a composite form.

In this work we study the setting where the potential U is the sum of a single smooth and a potentially large number of nonsmooth convex functions. In particular, we consider the problem

$$\text{Sample from } \mu^*(x) \propto \exp(-U(x)), \quad \text{where } U(x) := F(x) + \sum_{i=1}^n G_i(x), \quad (1)$$

*Also affiliated with Moscow Institute of Physics and Technology, Dolgoprudny, Russia.

Table 1: Complexity results obtained in Corollaries 2, 3 and 4 of our main result (Theorem 1).

F	Stepsize γ	Rate	Theorem
convex	$\mathcal{O}(\varepsilon)$	$\text{KL}(\mu_{\hat{x}^k} \mu^*) \leq \frac{1}{2\gamma(k+1)} W^2(\mu_{x^0}, \mu^*) + \mathcal{O}(\gamma)$	Cor 2
α -strongly convex	$\mathcal{O}(\varepsilon\alpha)$	$W^2(\mu_{x^k}, \mu^*) \leq (1 - \gamma\alpha)^k W^2(\mu_{x^0}, \mu^*) + \mathcal{O}(\frac{\gamma}{\alpha})$	Cor 3
α -strongly convex	$\mathcal{O}(\varepsilon\alpha)$	$\text{KL}(\mu_{\hat{x}^k} \mu^*) \leq \alpha(1 - \gamma\alpha)^{k+1} W^2(\mu_{x^0}, \mu^*) + \mathcal{O}(\gamma)$	Cor 4

where $F : \mathbb{R}^d \rightarrow \mathbb{R}$ is a smooth convex function and $G_1, \dots, G_n : \mathbb{R}^d \rightarrow \mathbb{R}$ are (possibly nonsmooth) convex functions. The additive model for U offers ample flexibility as typically there are multiple decompositions of U in the form (1).

2 Contributions

We now briefly comment some of the key contributions of this work.

◊ **A splitting technique for Langevin algorithm.** We propose a new variant of LA for solving (1), which we call *Stochastic Proximal Langevin Algorithm (SPLA)*. We assume that F and G can be written as expectations over some simpler functions $f(\cdot, \xi)$ and $g_i(\cdot, \xi)$

$$F(x) = \mathbb{E}_\xi(f(x, \xi)), \quad \text{and} \quad G_i(x) = \mathbb{E}_\xi(g_i(x, \xi)). \quad (2)$$

SPLA (see Algorithm 1 in Section 4) only requires accesses to the gradient of $f(\cdot, \xi)$ and to proximity operators of the functions $g_i(\cdot, \xi)$. SPLA can be seen as a Langevin version of the stochastic Passty algorithm [30, 36]. To the best of our knowledge, this is the first time a splitting technique that involves multiple (stochastic) proximity operators is used in a Langevin algorithm.

Remarks: Current forms of LA tackle problem (1) using stochastic subgradients [18]. If $n = 1$ and G_1 is proximal (i.e., the learner has access to the full proximity operator of G_1), it has recently been proposed to use proximity operators instead of (sub)gradients [20, 18], as it is done in the optimization literature [29, 2]. Indeed, in this case, the proximal stochastic gradient method is an efficient method to minimize U . If $n > 1$, and the functions G_i are proximal (but not U), the minimization of U is usually tackled using the *operator splitting* framework: the (stochastic) three-operator splitting [51, 17] or (stochastic) primal dual algorithms [13, 46, 9, 35]. These algorithms involve the computation of (stochastic) gradients and (full) proximity operators and enjoy numerical stability properties. However, proximity operators are sometimes difficult to implement. In this case, *stochastic proximity operators* are cheaper² than full proximity operators and numerically more stable than stochastic subgradients to handle nonsmooth terms [32, 31, 4, 5, 6] but also smooth [43] terms. In this paper, we bring together the advantages of operator splitting and stochastic proximity operators for sampling purposes.

◊ **Theory.** We perform a nonasymptotic convergence analysis of SPLA. Our main result, Theorem 1, gives a tractable recursion involving the Kullback-Leibler divergence and Wasserstein distance (when U is strongly convex) between μ^* and the distribution of certain samples generated by our method. We use this result to show that the KL divergence is lower than ε after $\mathcal{O}(1/\varepsilon^2)$ iterations if the constant stepsize $\gamma = \mathcal{O}(\varepsilon)$ is used (Corollary 2). Assuming F is α -strongly convex, we show that the Wasserstein distance and (resp. the KL divergence) decrease exponentially, up to an oscillation region of size $\mathcal{O}(\gamma/\alpha)$ (resp. $\mathcal{O}(\gamma)$) as shown in Corollary 3 (resp. Corollary 4). If we wish to push the Wasserstein distance below ε (resp. the KL divergence below $\alpha\varepsilon$), this could be achieved by setting $\gamma = \mathcal{O}(\varepsilon\alpha)$, and it would be sufficient to take $\mathcal{O}(1/\varepsilon \log 1/\varepsilon)$ iterations. These results are summarized in Table 1. The obtained convergence rates match the previous known results obtained in simpler settings [18]. Note that convergence rates of optimization methods involving multiple stochastic proximity operators haven't been established yet.

²See www.proximity-operator.net

Remarks: Our proof technique is inspired by [38], which is itself based on [18]. In [38], the authors consider the $n = 1$ case, and assume that the smooth function F is proximable. In [18], a proximal stochastic (sub)gradient Langevin algorithm is studied. In this paper, convergence rates are established by showing that the probability distributions of the iterates shadow some discretized gradient flow defined on a measure space. Hence, our work is a contribution to recent efforts to understand Langevin algorithm as an optimization algorithm in a space of probability measures [26, 49, 3].

◊ **Online setting.** In online settings, U is unknown but revealed across time. Our approach provides a reasonable algorithm for such situations, especially in cases when the information revealed about U is stationary in time. In particular, this includes online Bayesian learning with structured priors or nonsmooth log likelihood [50, 23, 40, 47]. In this context, the learner is required to sample from some posterior distribution μ^* that takes the form (1) where F, G_1, \dots, G_n are intractable. However, these functions can be cheaply sampled, or are revealed across time through i.i.d. streaming data.

◊ **Simulations.** We illustrate the promise of our approach numerically by performing experiments with SPLA. We first apply SPLA to a stochastic and nonsmooth toy model with a ground truth. Then, we consider the problem of *sampling from the posterior distribution in the Graph Trend Filtering* context [47]. For this nonsmooth large scale simulation problem, SPLA is performing better than the state of the art method that uses stochastic subgradients instead of stochastic proximity operators. Indeed, in the optimization literature [2], proximity operators are already known to be more stable than subgradients.

3 Technical Preliminaries

In this section, we recall certain notions from convex analysis and probability theory, which are keys to the developments in this paper, state our main assumptions, and introduce needed notations.

3.1 Subdifferential, minimal section and proximity operator

Given a convex function $g : \mathbb{R}^d \rightarrow \mathbb{R}$, its *subdifferential* at x , $\partial g(x)$, is the set

$$\partial g(x) := \{d \in \mathbb{R}^d : g(x) + \langle d, y - x \rangle \leq g(y)\}.$$

Since $\partial g(x)$ is a nonempty closed convex set [2], the projection of 0 onto $\partial g(x)$ —the least norm element in the set $\partial g(x)$ —is well defined, and we call this element $\nabla^0 g(x)$. The function $\nabla^0 g : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is called the *minimal section* of ∂g . The *proximity operator* associated with g is the mapping $\text{prox}_g : \mathbb{R}^d \rightarrow \mathbb{R}^d$ defined by

$$\text{prox}_g(x) := \arg \min_{y \in \mathbb{R}^d} \left\{ \frac{1}{2} \|x - y\|^2 + g(y) \right\}.$$

Due to its implicit definition, prox_g can be hard to evaluate.

3.2 Stochastic structure of F and G_i : integrability, smoothness and convexity

Here we detail the assumptions behind the stochastic structure (2) of the functions $F = \mathbb{E}_\xi(f(x, \xi))$ and $G_i = \mathbb{E}_\xi(g_i(x, \xi))$ defining the potential U . Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and denote \mathbb{E} the mathematical expectation and \mathbb{V} the variance. Consider ξ a random variable from Ω to another probability space (Ξ, \mathcal{G}) with distribution μ .

Assumption 1 (Integrability). The functions $f : \mathbb{R}^d \times \Xi \rightarrow \mathbb{R}^d$ and $g_i : \mathbb{R}^d \times \Xi \rightarrow \mathbb{R}^d, i = 1, \dots, n$, are μ -integrable for every $x \in \mathbb{R}^d$.

Furthermore, we will make the following convexity and smoothness assumptions.

Assumption 2 (Convexity and differentiability). The function $f(\cdot, s)$ is convex and differentiable for every $s \in \Xi$. The functions $g_i(\cdot, s)$ are convex for every $i \in \{1, 2, \dots, n\}$.

The gradient of $f(\cdot, s)$ is denoted $\nabla f(\cdot, s)$, the subdifferential of $g_i(\cdot, s)$ is denoted $\partial g_i(\cdot, s)$ and its minimal section is denoted $\nabla^0 g_i(\cdot, s)$. Under Assumption 2, it is known that F is convex and differentiable and that $\nabla F(x) = \mathbb{E}_\xi(\nabla f(x, \xi))$ [34]. Next, we assume that F is smooth and α -strongly convex. However, we allow $\alpha = 0$ if F is not strongly convex. We will only assume that $\alpha > 0$ in Corollaries 3 and 4.

Assumption 3 (Convexity and smoothness of F). The gradient of F is L -Lipschitz continuous, where $L \geq 0$. Moreover, F is α -strongly convex, where $\alpha \geq 0$.

Under Assumption 2, the second part of the above holds for $\alpha = 0$. Finally, we will introduce two noise conditions on the stochastic (sub)gradients of $f(\cdot, s)$ and $g_i(\cdot, s)$.

Assumption 4 (Bounded variance of $\nabla f(x, \cdot)$). There exists $\sigma_F \geq 0$, such that $\mathbb{V}_\xi(\|\nabla f(x, \xi)\|) \leq \sigma_F^2$ for every $x \in \mathbb{R}^d$.

Assumption 5 (Bounded second moment of $\nabla^0 g_i(x, \cdot)$). For every $i \in \{1, 2, \dots, n\}$, there exists $L_{G_i} \geq 0$ such that $\mathbb{E}_\xi(\|\nabla^0 g_i(x, \xi)\|^2) \leq L_{G_i}^2$ for every $x \in \mathbb{R}^d$.

Note that if $g_i(\cdot, s)$ is $\ell_i(s)$ -Lipschitz continuous for every $s \in \Xi$, and if $\ell_i(s)$ is μ -square integrable, then Assumption 5 holds.

3.3 KL divergence, entropy and potential energy

Recall from (1) that $U := F + \sum_{i=1}^n G_i$ and assume that $\int \exp(-U(x)) dx < \infty$. Our goal is to sample from the unique distribution μ^* over \mathbb{R}^d with density $\mu^*(x)$ (w.r.t. the Lebesgue measure denoted \mathcal{L}) proportional to $\exp(-U(x))$, for which we write $\mu^*(x) \propto \exp(-U(x))$. The closeness between the samples of our algorithm and the target distribution μ^* will be evaluated in terms of information theoretic and optimal transport theoretic quantities.

Let $\mathcal{B}(\mathbb{R}^d)$ be the Borel σ -field of \mathbb{R}^d . Given two nonnegative measures μ and ν on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$, we write $\mu \ll \nu$ if μ is absolutely continuous w.r.t. ν , and denote $\frac{d\mu}{d\nu}$ its density. The *Kullback-Leibler (KL) divergence* between μ and ν , $\text{KL}(\mu \mid \nu)$, quantifies the closeness between μ and ν . If $\mu \ll \nu$, then the KL divergence is defined by

$$\text{KL}(\mu \mid \nu) := \int \log \left(\frac{d\mu}{d\nu}(x) \right) d\mu(x),$$

and otherwise we set $\text{KL}(\mu \mid \nu) = +\infty$. Up to an additive constant, $\text{KL}(\cdot \mid \mu^*)$ can be seen as the sum of two terms [37]: the *entropy* $\mathcal{H}(\mu)$ and the *potential energy* $\mathcal{E}_U(\mu)$. The entropy of μ is given by $\mathcal{H}(\mu) := \text{KL}(\mu \mid \mathcal{L})$, and the potential energy of μ is defined by $\mathcal{E}_U(\mu) := \int U d\mu(x)$.

3.4 Wasserstein distance

Although the KL divergence is equal to zero if and only if $\mu = \nu$, it is not a mathematical distance (metric). The *Wasserstein distance*, defined below, metrizes the space $\mathcal{P}_2(\mathbb{R}^d)$ of probability measures over \mathbb{R}^d with a finite second moment. Consider $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$. A *transference plan* of (μ, ν) is a probability measure v over $(\mathbb{R}^d \times \mathbb{R}^d, \mathcal{B}(\mathbb{R}^d \times \mathbb{R}^d))$ with marginals μ, ν : for every $A \in \mathcal{B}(\mathbb{R}^d)$, $v(A \times \mathbb{R}^d) = \mu(A)$ and $v(\mathbb{R}^d \times A) = \nu(A)$. In particular, the product measure $\mu \otimes \nu$ is a transference plan. We denote $\Gamma(\mu, \nu)$ the set of transference plans. A *coupling* of (μ, ν) is a random variable (X, Y) over some probability space with values in $(\mathbb{R}^d \times \mathbb{R}^d, \mathcal{B}(\mathbb{R}^d \times \mathbb{R}^d))$ (i.e., X and Y are random variables with values in \mathbb{R}^d) such that the distribution of X is μ and the distribution of Y is ν . In other words, (X, Y) is a coupling of μ, ν if the distribution of (X, Y) is a transference plan of μ, ν . The Wasserstein distance of order 2 between μ and ν is defined by

$$W^2(\mu, \nu) := \inf \left\{ \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 dv(x, y), \quad v \in \Gamma(\mu, \nu) \right\}.$$

One can see that $W^2(\mu, \nu) = \inf \mathbb{E}(\|X - Y\|^2)$, where the inf is taken over all couplings (X, Y) of μ, ν defined on some probability space with expectation \mathbb{E} .

4 The SPLA Algorithm and its Convergence Rates

4.1 The algorithm

To solve the sampling problem (1), our Stochastic Proximal Langevin Algorithm (SPLA) generates a sequence of random variables $(x^k)_{k \geq 0}$ from $(\Omega, \mathcal{F}, \mathbb{P})$ to $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ defined as follows

$$\begin{aligned} z^k &= x^k - \gamma \nabla f(x^k, \xi^k) \\ y_0^k &= z^k + \sqrt{2\gamma} W^k \\ y_i^k &= \text{prox}_{\gamma g_i(\cdot, \xi^k)}(y_{i-1}^k) \quad \text{for } i = 1, \dots, n \\ x^{k+1} &= y_n^k, \end{aligned}$$

where $(W^k)_{k \geq 0}$ is a sequence of i.i.d. standard Gaussian random variables, $(\xi^k)_{k \geq 0}$ is a sequence of i.i.d. copies of ξ and $\gamma > 0$ is a positive step size. Our SPLA method is formalized as Algorithm 1; its steps are explained therein.

Algorithm 1 Stochastic Proximal Langevin Algorithm (SPLA)

Initialize: $x^0 \in \mathbb{R}^d$
for $k = 0, 1, 2, \dots$ **do**
 Sample random ξ^k ▷ used for stoch. approximation: $F \approx f(\cdot, \xi^k)$ and $G_i \approx g_i(\cdot, \xi^k)$
 $z^k = x^k - \gamma \nabla f(x^k, \xi^k)$ ▷ a stochastic gradient descent step in F
 Sample random W^k ▷ a standard Gaussian vector in \mathbb{R}^d
 $y_0^k = z^k + \sqrt{2\gamma} W^k$ ▷ a Langevin step w.r.t. F
 for $i = 1, \dots, n$ **do**
 $y_i^k = \text{prox}_{\gamma g_i(\cdot, \xi^k)}(y_{i-1}^k)$ ▷ prox step to handle the term $G_i(\cdot) = \mathbb{E}_\xi g_i(\cdot, \xi)$
 end for
 $x^{k+1} = y_n^k$ ▷ the final SPLA step, accounting for F and G_1, G_2, \dots, G_n
end for

4.2 Main theorem

We now state our main results in terms of Kullback-Leibler divergence and Wasserstein distance. We denote μ_x the distribution of every random variable x defined on $(\Omega, \mathcal{F}, \mathbb{P})$.

Theorem 1. Let Assumptions 1–5 hold and assume that $\gamma \leq 1/L$. There exists $C \geq 0$ such that,

$$2\gamma \text{KL}(\mu_{y_0^k} \mid \mu^*) \leq (1 - \gamma\alpha) W^2(\mu_{x^k}, \mu^*) - W^2(\mu_{x^{k+1}}, \mu^*) + \gamma^2(2\sigma_F^2 + 2Ld + C). \quad (3)$$

The constant C can be expressed as a linear combination of $L_{G_1}^2, \dots, L_{G_n}^2$ with integer coefficients. Moreover, if $n = 2$, then $C := 2(L_{G_1}^2 + L_{G_2}^2)$. More generally, if for every $i \in \{2, \dots, n\}$, $g_i(\cdot, \xi)$ admits almost surely the representation $g_i(\cdot, \xi) = \tilde{g}_i(\cdot, \xi_i)$ where ξ_2, \dots, ξ_n are independent random variables, then $C := n \sum_{i=1}^n L_{G_i}^2$.

Proof. A full proof can be found in the Supplementary material. We only sketch the main steps here. For every μ -integrable function $g : \mathbb{R}^d \rightarrow \mathbb{R}$, we denote $\mathcal{E}_g(\mu) = \int g d\mu$. Moreover, we denote $\mathcal{F} = \mathcal{E}_U + \mathcal{H}$. First, using [18, Lemma 1], $\mu^* \in \mathcal{P}_2(\mathbb{R}^d)$, $\mathcal{E}_U(\mu^*), \mathcal{H}(\mu^*) < \infty$ and if $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, then

$$\text{KL}(\mu \mid \mu^*) = \mathcal{E}_U(\mu) + \mathcal{H}(\mu) - (\mathcal{E}_U(\mu^*) + \mathcal{H}(\mu^*)) = \mathcal{F}(\mu) - \mathcal{F}(\mu^*),$$

provided that $\mathcal{E}_U(\mu) < \infty$. Then, we decompose $\mathcal{E}_U(\mu) = \mathcal{E}_F(\mu) + \mathcal{E}_G(\mu)$ where $G = \sum_{i=1}^n G_i$. Using [18] again, we can establish the inequality

$$2\gamma \left[\mathcal{H}(\mu_{y_0^k}) - \mathcal{H}(\mu^*) \right] \leq W^2(\mu_{z^k}, \mu^*) - W^2(\mu_{y_0^k}, \mu^*). \quad (4)$$

Then, if $\gamma \leq 1/L$ we obtain, for every random variable a with distribution μ^* ,

$$\mathbb{E} \left[\|z^k - a\|^2 \right] \leq (1 - \gamma\alpha) \mathbb{E} \left[\|x^k - a\|^2 \right] + 2\gamma [\mathcal{E}_F(\mu^*) - \mathcal{E}_F(\mu_{z^k})] + 2\gamma^2 \sigma_F^2, \quad (5)$$

using standard computations regarding the stochastic gradient descent algorithm. Using the smoothness of F and the definition of the Wasserstein distance, this implies

$$2\gamma \left[\mathcal{E}_F(\mu_{y_0^k}) - \mathcal{E}_F(\mu^*) \right] \leq (1 - \gamma\alpha)W^2(\mu_{x^k}, \mu^*) - W^2(\mu_{z^k}, \mu^*) + \gamma^2(2\sigma_F^2 + 2Ld).$$

It remains to establish $2\gamma \left[\mathcal{E}_G(\mu_{y_0^k}) - \mathcal{E}_G(\mu^*) \right] \leq W^2(\mu_{y_0^k}, \mu^*) - W^2(\mu_{x^{k+1}}, \mu^*) + \gamma^2 C$, which is the main technical challenge of the proof. This is done using the frameworks of Yosida approximation of random subdifferentials and Moreau regularizations of random convex functions [2]. Equation (3) is obtained by summing the obtained inequalities. \square

4.3 Link with Wasserstein Gradient Flows

Equation (3) is reminiscent of the fact that the SPLA shadows the gradient flow of $\text{KL}(\cdot \mid \mu^*)$ in the metric space $(\mathcal{P}_2(\mathbb{R}^d), W)$. To see this, first consider the gradient flow associated to F . By definition, it is the flow of the differential equation [7]

$$\frac{d}{dt}x(t) = -\nabla F(x(t)), \quad t > 0. \quad (6)$$

The function x can alternatively be defined as a solution of the variational inequalities

$$2(F(x(t)) - F(a)) \leq -\frac{d}{dt}\|x(t) - a\|^2, \quad t > 0, \quad \forall a \in \mathbb{R}^d. \quad (7)$$

The iterates $(u^k)_{k \geq 0}$ of the stochastic gradient descent (SGD) algorithm applied to F can be seen as a (noisy) Euler discretization of (6) with a step size $\gamma > 0$. This idea has been used successfully in the stochastic approximation literature [33, 24]. This analogy goes further since a fundamental inequality used to analyze SGD applied to F is ([27])

$$2\gamma \mathbb{E}(F(u^{k+1}) - F(a)) \leq \mathbb{E}\|u^k - a\|^2 - \mathbb{E}\|u^{k+1} - a\|^2 + \gamma^2 K, \quad k \geq 0,$$

where $K \geq 0$ is some constant, which can be seen as a discrete counterpart of (7). Note that this inequality is similar to (5) that is used in the proof of Theorem 1.

In the optimal transport theory, the point of view of (7) is used to define the gradient flow of a (geodesically) convex function \mathcal{F} defined on $\mathcal{P}_2(\mathbb{R}^d)$ (see [37] or [1, Page 280]). Indeed, the gradient flow $(\nu_t)_{t \geq 0}$ of \mathcal{F} in the space $(\mathcal{P}_2(\mathbb{R}^d), W)$ satisfies for every $t > 0, \mu \in \mathcal{P}_2(\mathbb{R}^d)$,

$$2(\mathcal{F}(\nu_t) - \mathcal{F}(\mu)) \leq -\frac{d}{dt}W^2(\nu_t, \mu), \quad (8)$$

which can be seen as a continuous time counterpart of Equation (3) by setting $\mathcal{F} = \text{KL}(\cdot \mid \mu^*)$. Furthermore, Equation (4) in the proof of Theorem 1 is also related to (8). It is obtained by applying Equation (8) with $\mathcal{F} = \mathcal{H}$ and $\nu_0 = \mu_{z^k}$ (see e.g [18, Lemma 5]).

4.4 Explicit convergence rates for convex and strongly convex F

Corollaries 2, 3 and 4 below are obtained by unrolling the recursion provided by Theorem 1. The results are summarized in Table 1.

Corollary 2 (Convex F). Consider a sequence of independent random variables $(j_k)_{k \geq 0}$ such that $(j_k)_{k \geq 0}$ is independent of $(W^k)_k$ and $(\xi^k)_k$, and the distribution of j_k is uniform over $\{0, \dots, k\}$. Denote $\hat{x}^k = y_0^{j_k}$. If $\gamma \leq 1/L$, then,

$$\text{KL}(\mu_{\hat{x}^k} \mid \mu^*) \leq \frac{1}{2\gamma(k+1)}W^2(\mu_{x^0}, \mu^*) + \frac{\gamma}{2}(2\sigma_F^2 + 2Ld + C).$$

Hence, given any $\varepsilon > 0$, choosing stepsize $\gamma = \min \left\{ \frac{1}{L}, \frac{\varepsilon}{2\sigma_F^2 + 2Ld + C} \right\}$ and a number of iterations

$$k + 1 \geq \max \left\{ \frac{L}{\varepsilon}, \frac{2\sigma_F^2 + 2Ld + C}{\varepsilon^2} \right\} W^2(\mu_{x^0}, \mu^*),$$

implies $\text{KL}(\mu_{\hat{x}^k} \mid \mu^*) \leq \varepsilon$.

Corollary 3 (Strongly convex F). If $\alpha > 0$ and $\gamma \leq 1/L$, then,

$$W^2(\mu_{x^k}, \mu^*) \leq (1 - \gamma\alpha)^k W^2(\mu_{x^0}, \mu^*) + \frac{\gamma(2\sigma_F^2 + 2Ld + C)}{\alpha}.$$

Hence, given any $\varepsilon > 0$, choosing stepsize $\gamma = \min \left\{ \frac{1}{L}, \frac{\varepsilon\alpha}{2(2\sigma_F^2 + 2Ld + C)} \right\}$ and a number of iterations

$$k \geq \max \left\{ \frac{L}{\alpha}, \frac{2(2\sigma_F^2 + 2Ld + C)}{\varepsilon\alpha^2} \right\} \log \left(\frac{2W^2(\mu_{x^0}, \mu^*)}{\varepsilon} \right),$$

implies $W^2(\mu_{x^k}, \mu^*) \leq \varepsilon$.

Corollary 4 (Strongly convex F). Consider a sequence of independent random variables $(j_k)_{k \geq 0}$ such that $(j_k)_k$ is independent of $(W^k)_k$ and $(\xi^k)_k$. Assume that the distribution of j_k is geometric over $\{0, \dots, k\}$:

$$\mathbb{P}(j_k = r) \propto (1 - \gamma\alpha)^{-r}.$$

Denote $\tilde{x}^k = x^{j_k}$. If $\alpha > 0$ and $\gamma \leq 1/L$, then,

$$\text{KL}(\mu_{\tilde{x}^k} \mid \mu^*) \leq \frac{\alpha W^2(\mu_{x^0}, \mu^*)}{2} \cdot \frac{(1 - \gamma\alpha)^{k+1}}{1 - (1 - \gamma\alpha)^{k+1}} + \frac{\gamma(2\sigma_F^2 + 2Ld + C)}{2}.$$

Hence, given any $\varepsilon > 0$, choosing stepsize $\gamma = \min \left\{ \frac{1}{L}, \frac{\varepsilon\alpha}{2\sigma_F^2 + 2Ld + C} \right\}$ and a number of iterations

$$k \geq \max \left\{ \frac{L}{\alpha}, \frac{2\sigma_F^2 + 2Ld + C}{\varepsilon\alpha^2} \right\} \log \left(2 \max \left\{ 1, \frac{W^2(\mu_{x^0}, \mu^*)}{\varepsilon} \right\} \right),$$

implies $\text{KL}(\mu_{\tilde{x}^k} \mid \mu^*) \leq \alpha\varepsilon$.

We can compare these bounds with the one of [18]. First, in the particular case $n = 1$ and $g_1(\cdot, s) \equiv G_1$, SPLA boils down to the algorithm of [18, Section 4.2], Corollary 2 matches exactly [18, Corollary 18] and Corollary 3 matches [18, Corollary 22]. To our knowledge, Corollary 4 has no counterpart in the litterature. We now focus on the case $F \equiv 0$ and $n = 1$ of SPLA, as it concentrates the innovations of our paper. In this case, $L = 0$ and $\sigma_F = 0$. Compared to the Stochastic Subgradient Langevin Algorithm (SSLA) [18, Section 4.1], Corollary 2 matches with [18, Corollary 14].

5 Numerical experiments

5.1 Simulations using a ground truth

We first concentrate on the case $F \equiv 0$ and $n = 1$. Let $U = |x| = \mathbb{E}_\xi(|x| + x\xi)$ ($g_1(x, s) = |x| + xs$), where ξ is standard Gaussian. The target $\mu^* \propto \exp(-U)$ is a standard Laplace distribution in \mathbb{R} . In this case, $L = \alpha = \sigma_F = 0$ and $C = L_{G_1}^2 = 2$. We shall illustrate the bound on $\text{KL}(\mu_{\hat{x}^k} \mid \mu^*)$ (Corollary 2 for SPLA and [18, Corollary 14] for SSLA) for both algorithms using histograms. Note that the distribution $\mu_{\hat{x}^k}$ of \hat{x}^k is a (deterministic) mixture of the μ_{x^j} : $\mu_{\hat{x}^k} = \frac{1}{k} \sum_{j=1}^k \mu_{x^j}$. Using Pinsker inequality, we can bound the total variation distance between $\mu_{\hat{x}^k}$ and μ^* from the bound on KL, and illustrate this by histograms. In Figure 1, we take $\gamma = 10$ and do 10^5 iterations of both algorithms. Note that here the complexity of one iteration of SPLA or SSLA is the same. One can




Figure 1: Comparison between histograms of SPLA and SSLA and the true density $0.5 \exp(-|x|)$.

see that SPLA enjoys the well known advantages of stochastic proximal methods [43]: precision, numerical stability (less outliers), and robustness to step size.

Algorithm 2 SPLA for the Graph Trend Filtering

Initialize: $x^0 \in \mathbb{R}^V$
for $k = 0, 1, 2, \dots$ **do**
 $z^k = x^k - \frac{\gamma}{\sigma^2}(x^k - Y)$
 Sample random W^k ▷ standard Gaussian vector in \mathbb{R}^V
 $y_0^k = z^k + \sqrt{2\gamma}W^k$
for $i = 1, \dots, n$ **do**
 Sample uniform random edges e_i
 $y_i^k = \text{prox}_{\gamma g_{e_i}}(y_{i-1}^k)$
end for
 $x^{k+1} = y_n^k$
end for

5.2 Application to Trend Filtering on Graphs

In this section we consider the following Bayesian point of view of *trend filtering on graphs* [42]. Consider a finite undirected graph $G = (V, E)$, where V is the set of vertices and E is the set of edges. Denote d the cardinality of V and $|E|$ the cardinality of E . A realization of a random vector $Y \in \mathbb{R}^V$ is observed. In a Bayesian framework, the distribution of Y is parametrized by a vector $X \in \mathbb{R}^V$ which is itself random and whose distribution p is proportional to $\exp(-\lambda \text{TV}(x, G))$, where $\lambda > 0$ is a scaling parameter and where for every $x \in \mathbb{R}^V$

$$\text{TV}(x, G) = \sum_{i,j \in V, \{i,j\} \in E} |x(i) - x(j)|,$$

is the Total Variation regularization over G . The goal is to learn X after an observation of Y . The paper [47] consider the case where the distribution of Y given X (a.k.a the likelihood) is proportional to $\exp(-\frac{1}{2\sigma^2}\|X - y\|^2)$, where $\sigma \geq 0$ is another scaling parameter. In other words, the distribution of Y given X is $N(X, \sigma^2 I)$, a normal distribution centered at X with variance $\sigma^2 I$ (where I is the $d \times d$ identity matrix). Denoting

$$\pi(x | y) \propto \exp(-U(x)), \quad U(x) = \frac{1}{2\sigma^2}\|x - y\|^2 + \lambda \text{TV}(x, G),$$

the posterior distribution of X given Y , the maximum a posteriori estimator in this Bayesian framework is called the Graph Trend Filtering estimate [47]. It can be written

$$x^* = \arg \max_{x \in \mathbb{R}^V} \pi(x | Y) = \arg \min_{x \in \mathbb{R}^V} \frac{1}{2\sigma^2}\|x - Y\|^2 + \lambda \text{TV}(x, G).$$

Although maximum a posteriori estimators carry some information, they are not able to capture uncertainty in the learned parameters. Samples a posteriori provide a better understanding of the posterior distribution and allow to compute other Bayesian estimates such as confidence intervals. This allows to avoid overfitting among other things. In our context, sampling a posteriori would require to sample from the target distribution $\mu^*(x) = \pi(x | Y)$.

In the case where G is a 2D grid (which can be identified with an image), the proximity operator of $\text{TV}(\cdot, G)$ can be computed using a subroutine [8] and the proximal stochastic gradient Langevin algorithm can be used to sample from $\pi(\cdot | Y)$ [20, 18]. However, on a large/general graph, the proximity operator of $\text{TV}(\cdot, G)$ is hard to evaluate [41, 36]. Since $\text{TV}(\cdot, G)$ is written as a sum, we shall rather select a batch of random edges and compute the proximity operators over these randomly chosen edges. More precisely, we write the potential U defining $\pi(x | Y)$ in the form (1) by setting

$$U(x) = F(x) + \sum_{i=1}^n G_i(x), \quad F(x) = \frac{1}{2\sigma^2}\|x - Y\|^2, \quad G_i(x) = \lambda \frac{|E|}{n} \mathbb{E}_{e_i} (|x(v_i) - x(w_i)|),$$

where for every $i \in \{1, \dots, n\}$, $e_i = \{v_i, w_i\} \in E$ is an uniform random edge and the e_i are independent. For every edge $e = \{v, w\} \in E$, (where v, w are vertices) denote $g_e(x) = \lambda \frac{|E|}{n} |x(v) - x(w)|$ and note that $G_i(x) = \mathbb{E}_{e_i}(g_{e_i}(x))$. The parameter n can be seen as a batch parameter: $\sum_{i=1}^n g_{e_i}(x)$ is an unbiased approximation of $\text{TV}(x, G)$. Also note that we set $f(\cdot, s) \equiv F$. The SPLA applied to sample from $\pi(\cdot | Y)$ is presented as Algorithm 2. In our simulations, the SPLA is compared to two different versions of the Langevin algorithm. In the

Stochastic Subgradient Langevin Algorithm (SSLA) [18], stochastic subgradients of g_{e_i} are used instead of stochastic proximity operators. In the Proximal Langevin Algorithm (ProxLA) [18], the full proximity operator of $\sum_{i=1}^n G_i$ is computed using a subroutine. As mentioned in [36, 47], we use the gradient algorithm for the dual problem. The plots in Figure 2 provide simulations of the algorithms on our machine (using one thread of a 2,800 MHz CPU and 256GB RAM). Additional numerical experiments are available in the Appendix. Four real life graphs from the dataset [25] are considered : the Facebook graph (4,039 nodes and 88,234 edges, extracted from the Facebook social network), the Youtube graph (1,134,890 nodes and 2,987,624 edges, extracted from the social network included in the Youtube website), the Amazon graph (the 334,863 nodes represent products linked by and 925,872 edges) and the DBLP graph (a co-authorship network of 317,080 nodes and 1,049,866 edges). On the larger graphs, we only simulate SPLA and SSLA since the computation of a full proximity operator becomes prohibitive. Numerical experiments over the Amazon and the DBLP graphs are available in the Supplementary material.

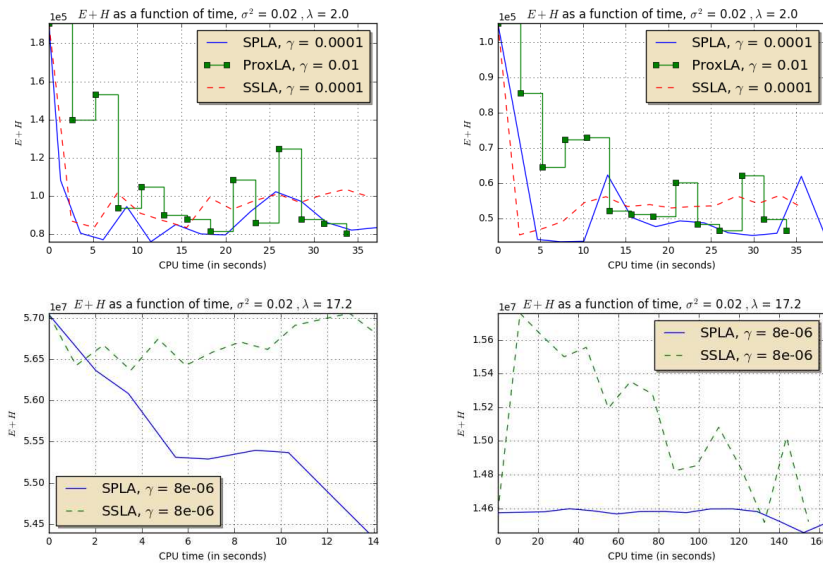


Figure 2: **Top row:** The functional $\mathcal{F} = \mathcal{H} + \mathcal{E}_U$ as a function of CPU time for the three algorithms over the Facebook graph. Left: $Y \sim N(0, I)$. Right: $Y \sim N(0, I)$ and then half of the coordinates of Y are put to zero. **Bottom row:** The functional $\mathcal{F} = \mathcal{H} + \mathcal{E}_U$ as a function of CPU time for the two algorithms over the Youtube graph. Left: $Y \sim N(0, I)$. Right: $Y \sim N(0, I)$ and then half of the coordinates of Y are put to zero.

In our simulations, we represent the functional $\mathcal{F} = \mathcal{H} + \mathcal{E}_U$ as a function of CPU time while running the algorithms. The parameters λ and σ are chosen such that the log likelihood term and the Total Variation regularization term have the same weight. The functionals \mathcal{H} and \mathcal{E}_U are estimated using five random realizations of each iterate \hat{x}^k (\mathcal{H} is estimated using a kernel density estimator). The batch parameter n is equal to 400. We consider cases where Y has a standard gaussian distribution and cases where half of the components of Y are standard gaussians and half are equal to zero (this correspond to the graph inpainting task [11]). SPLA and SSLA are always simulated with the same step size.

As expected, the numerical experiments show the advantage of using stochastic proximity operators instead of stochastic subgradients. It is a standard fact that proximity operators are better than subgradients to tackle ℓ^1 -norm terms [2]. Our figures show that stochastic proximity operators are numerically more stable than alternatives [43]. Our figures also show the advantage of stochastic methods (SSLA or SPLA) over deterministic ones for large scale problems. The SSLA and the SPLA provide iterates about one hundred times more frequently than ProxLA, and are faster in the first iterations.

References

- [1] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient Flows: In Metric Spaces and in the Space of Probability Measures*. Springer Science & Business Media, 2008.
- [2] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. CMS Books in Mathematics/Ouvrages de Mathématiques de la SMC. Springer, New York, 2011.
- [3] E. Bernton. Langevin Monte Carlo and JKO splitting. *arXiv preprint arXiv:1802.08671*, 2018.
- [4] P. Bianchi. Ergodic convergence of a stochastic proximal point algorithm. *SIAM Journal on Optimization*, 26(4):2235–2260, 2016.
- [5] P. Bianchi and W. Hachem. Dynamical behavior of a stochastic Forward-Backward algorithm using random monotone operators. *Journal of Optimization Theory and Applications*, 171(1):90–120, 2016.
- [6] P. Bianchi, W. Hachem, and A. Salim. A constant step Forward-Backward algorithm involving random maximal monotone operators. *Journal of Convex Analysis*, 26(2):397–436, 2019.
- [7] H. Brézis. *Opérateurs maximaux monotones et semi-groupes de contractions dans les espaces de Hilbert*. North-Holland mathematics studies. Elsevier Science, Burlington, MA, 1973.
- [8] A. Chambolle, V. Caselles, D. Cremers, M. Novaga, and T. Pock. An introduction to total variation for image analysis. *Theoretical foundations and numerical methods for sparse recovery*, 9:263–340, 2010.
- [9] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.
- [10] N. S Chatterji, N. Flammarion, Y.-A. Ma, P. L Bartlett, and M. I Jordan. On the theory of variance reduction for stochastic gradient Monte Carlo. *arXiv preprint arXiv:1802.05431*, 2018.
- [11] S. Chen, A. Sandryhaila, G. Lederman, Z. Wang, J. MF Moura, P. Rizzo, J. Bielik, J. H Garrett, and J. Kovačević. Signal inpainting on graphs via total variation minimization. In *International Conference on Acoustics, Speech, and Signal Processing*, pages 8267–8271, 2014.
- [12] X. Cheng, N. S Chatterji, P. L Bartlett, and M. I Jordan. Underdamped Langevin MCMC: A non-asymptotic analysis. *arXiv preprint arXiv:1707.03663*, 2017.
- [13] L. Condat. A primal-dual splitting method for convex optimization involving Lipschitzian, proximable and linear composite terms. *Journal of Optimization Theory and Applications*, 158(2):460–479, 2013.
- [14] A. S Dalalyan. Theoretical guarantees for approximate sampling from smooth and log-concave densities. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(3):651–676, 2017.
- [15] A. S Dalalyan and A. Karagulyan. User-friendly guarantees for the langevin monte carlo with inaccurate gradient. *Stochastic Processes and their Applications*, 2019.
- [16] A. S Dalalyan and L. Riou-Durand. On sampling from a log-concave density using kinetic Langevin diffusions. *arXiv preprint arXiv:1807.09382*, 2018.
- [17] D. Davis and W. Yin. A three-operator splitting scheme and its optimization applications. *Set-Valued and Variational Analysis*, 25(4):829–858, 2017.
- [18] A. Durmus, S. Majewski, and B. Miasojedow. Analysis of Langevin Monte Carlo via convex optimization. *Journal of Machine Learning Research*, 20(73):1–46, 2019.
- [19] A. Durmus and E. Moulines. Nonasymptotic convergence analysis for the unadjusted Langevin algorithm. *The Annals of Applied Probability*, 27(3):1551–1587, 2017.

- [20] 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*, 11(1):473–506, 2018.
- [21] E. Hazan, K. Y Levy, and S. Shalev-Shwartz. On graduated optimization for stochastic non-convex problems. In *International conference on machine learning*, pages 1833–1841, 2016.
- [22] Y.-P. Hsieh, A. Kavis, P. Rolland, and V. Cevher. Mirrored Langevin dynamics. In *Advances in Neural Information Processing Systems*, pages 2878–2887, 2018.
- [23] S. Kotz, T. Kozubowski, and K. Podgorski. *The Laplace distribution and generalizations: a revisit with applications to communications, economics, engineering, and finance*. Springer Science & Business Media, 2012.
- [24] H. J. Kushner and G. G. Yin. *Stochastic approximation and recursive algorithms and applications*, volume 35 of *Applications of Mathematics (New York)*. Springer-Verlag, New York, second edition, 2003. Stochastic Modelling and Applied Probability.
- [25] J. Leskovec and A. Krevl. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>, June 2014.
- [26] Y.-A. Ma, N. Chatterji, X. Cheng, N. Flammarion, P. L Bartlett, and M. I Jordan. Is there an analog of Nesterov acceleration for MCMC? *arXiv preprint arXiv:1902.00996*, 2019.
- [27] E. Moulines and F. Bach. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In *Advances in Neural Information Processing Systems*, pages 451–459, 2011.
- [28] R. M Neal. MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2(11):2, 2011.
- [29] N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends® in Optimization*, 1(3):127–239, 2014.
- [30] G. B. Passty. Ergodic convergence to a zero of the sum of monotone operators in Hilbert space. *Journal of Mathematical Analysis and Applications*, 72(2):383–390, 1979.
- [31] A. Patrascu and I. Necoara. Nonasymptotic convergence of stochastic proximal point algorithms for constrained convex optimization. *Journal of Machine Learning Research*, 18:1–42, 2018.
- [32] P. Richtárik and M. Takáč. Stochastic reformulations of linear systems: algorithms and convergence theory. *arXiv:1706.01108*, 2017.
- [33] H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3):400–407, 1951.
- [34] R. T. Rockafellar and R. J.-B. Wets. On the interchange of subdifferentiation and conditional expectations for convex functionals. *Stochastics*, 7(3):173–182, 1982.
- [35] L. Rosasco, S. Villa, and B. C. Vũ. Stochastic inertial primal-dual algorithms. *arXiv preprint arXiv:1507.00852*, 2015.
- [36] A. Salim, P. Bianchi, and W. Hachem. Snake: a stochastic proximal gradient algorithm for regularized problems over large graphs. *IEEE Transactions on Automatic Control*, 2019.
- [37] F. Santambrogio. {Euclidean, metric, and Wasserstein} gradient flows: an overview. *Bulletin of Mathematical Sciences*, 7(1):87–154, 2017.
- [38] S. Schechtman, A. Salim, and P. Bianchi. Passty Langevin. In *CAp*, 2019.
- [39] U. Şimşekli. Fractional Langevin Monte Carlo: Exploring Lévy driven stochastic differential equations for Markov Chain Monte Carlo. In *International Conference on Machine Learning*, pages 3200–3209, 2017.

- [40] P. Sollich. Bayesian methods for support vector machines: Evidence and predictive class probabilities. *Machine learning*, 46(1-3):21–52, 2002.
- [41] W. Tansey and J. G Scott. A fast and flexible algorithm for the graph-fused lasso. *arXiv preprint arXiv:1505.06475*, 2015.
- [42] R. J. Tibshirani. Adaptive piecewise polynomial estimation via trend filtering. *The Annals of Statistics*, 42(1):285–323, 2014.
- [43] P. Toulis, T. Horel, and E. M Airoidi. Stable Robbins-Monro approximations through stochastic proximal updates. *arXiv preprint arXiv:1510.00967*, 2015.
- [44] T. Van Erven and P. Harremoës. Rényi divergence and Kullback-Leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.
- [45] C. Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.
- [46] B. C. Vũ. A splitting algorithm for dual monotone inclusions involving cocoercive operators. *Advances in Applied Mathematics*, 38(3):667–681, 2013.
- [47] Y.-X. Wang, J. Sharpnack, A. J Smola, and R. J Tibshirani. Trend filtering on graphs. *The Journal of Machine Learning Research*, 17(1):3651–3691, 2016.
- [48] M. Welling and Y. W Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *International Conference on Machine Learning*, pages 681–688, 2011.
- [49] A. Wibisono. Sampling as optimization in the space of measures: The Langevin dynamics as a composite optimization problem. *arXiv preprint arXiv:1802.08089*, 2018.
- [50] X. Xu and M. Ghosh. Bayesian variable selection and estimation for group lasso. *Bayesian Analysis*, 10(4):909–936, 2015.
- [51] A. Yurtsever, B. C. Vu, and V. Cevher. Stochastic three-composite convex minimization. In *Advances in Neural Information Processing Systems*, pages 4329–4337, 2016.
- [52] D. Zou, P. Xu, and Q. Gu. Subsampled stochastic variance-reduced gradient langevin dynamics. In *International Conference on Uncertainty in Artificial Intelligence*, 2018.
- [53] D. Zou, P. Xu, and Q. Gu. Sampling from non-log-concave distributions via variance-reduced gradient langevin dynamics. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2936–2945, 2019.