
Fractional Langevin Monte Carlo: Exploring Lévy Driven Stochastic Differential Equations for Markov Chain Monte Carlo

Umut Şimşekli¹

Abstract

Along with the recent advances in scalable Markov Chain Monte Carlo methods, sampling techniques that are based on Langevin diffusions have started receiving increasing attention. These so called Langevin Monte Carlo (LMC) methods are based on diffusions driven by a Brownian motion, which gives rise to Gaussian proposal distributions in the resulting algorithms. Even though these approaches have proven successful in many applications, their performance can be limited by the light-tailed nature of the Gaussian proposals. In this study, we extend classical LMC and develop a novel Fractional LMC (FLMC) framework that is based on a family of heavy-tailed distributions, called α -stable Lévy distributions. As opposed to classical approaches, the proposed approach can possess large jumps while targeting the correct distribution, which would be beneficial for efficient exploration of the state space. We develop novel computational methods that can scale up to large-scale problems and we provide formal convergence analysis of the proposed scheme. Our experiments support our theory: FLMC can provide superior performance in multi-modal settings, improved convergence rates, and robustness to algorithm parameters.

1. Introduction

Markov Chain Monte Carlo (MCMC) techniques that are based on continuous diffusions have become increasingly popular due to their success in large-scale Bayesian machine learning. In these techniques, the goal is to generate samples from a *target distribution* π , by forming a contin-

uous diffusion which has π as a stationary distribution. In practice, π is usually known up to a normalization constant, i.e. $\pi(X) \propto \phi(X) = \exp(-U(X))$ for $X \in \mathbb{R}^D$, where ϕ is called the *unnormalized target density* and U is called the *potential energy function*.

Originated in statistical physics (Rossky et al., 1978), Langevin Monte Carlo (LMC) is constructed upon the Langevin diffusion that is defined by the following stochastic differential equation (SDE) (Roberts & Stramer, 2002):

$$dX_t = -\nabla U(X_t)dt + \sqrt{2}dB_t, \quad (1)$$

where B_t denotes the standard D -dimensional Brownian motion. Under certain regularity conditions on U , the solution process $(X_t)_{t \geq 0}$ can be shown to be ergodic with π (Roberts & Stramer, 2002), which allows us to generate samples from π by simulating the continuous-time process (1) in discrete-time. This approach paves the way for the celebrated Unadjusted Langevin Algorithm (ULA) (Roberts & Stramer, 2002), that is given as follows:

$$\bar{X}_{n+1} = \bar{X}_n - \eta_{n+1} \nabla U(\bar{X}_n) + \sqrt{2\eta_{n+1}} \Delta B_{n+1}, \quad (2)$$

where n denotes the iterations, $(\eta_n)_n$ is a sequence of step-sizes, and $(\Delta B_n)_n$ is a sequence of independent and identically-distributed (i.i.d.) standard Gaussian random variables. Convergence properties of ULA have been studied in (Lamberton & Pages, 2003).

In a statistical physics context, X_t often represents the position of a particle (at time t) that is under the influence of a random force. In this case, the Langevin equation (1) is motivated by the hypothesis that this random force is the sum of many i.i.d. random ‘pulses’, whose variance is assumed to be finite (Yanovsky et al., 2000). Then, by the central limit theorem (CLT), the sum of these pulses converges to a Gaussian random variable, which justify the choice of the Brownian motion in the Langevin equation (1).

A natural question arises if we relax the finite variance assumption and allow the random pulses to have infinite variance. In such circumstances, the ‘usual’ CLT would not hold; however, one can still show that the sum of these pulses converges to a broader class of heavy-tailed distributions called α -stable (or Lévy-stable) distributions (Lévy,

¹LTCI, Télécom ParisTech, Université Paris-Saclay, 75013, Paris, France. Correspondence to: Umut Şimşekli <umut.simsekli@telecom-paristech.fr>.




Figure 1. Top: probability density functions of $\mathcal{S}\alpha\mathcal{S}$. Bottom: illustration of α -stable Lévy motions.

1937). Since the law of the random force is non-Gaussian in this case, the Brownian motion would not be appropriate in (1) and it needs to be replaced with the α -stable Lévy motion, which will be described in Section 2.

As opposed to the Brownian motion, which is almost surely continuous, the Lévy motion can contain discontinuities that are often referred to as ‘jumps’. Due to these jumps, the SDEs that are driven by Lévy motions are also called *anomalous diffusions*. It has been noticed that this heavy-tailed nature of the Lévy processes can be more appropriate for modeling natural phenomena that might incur large variations; a situation often encountered in statistical physics (Eliazar & Klafter, 2003), finance (Mandelbrot, 2013), and signal processing (Kuruoglu, 1999).

Despite the fact that Lévy-driven SDEs have been studied in more general Monte Carlo contexts (e.g. for financial simulations) (Konakov & Menozzi, 2011; Mikulevičius & Zhang, 2011), surprisingly, their use in MCMC has been left widely unexplored. In the statistical physics literature, Ditlevsen (1999) considered a Lévy-driven SDE with a double-well potential and investigated its waiting-times. In a similar context, Eliazar & Klafter (2003) developed an approximate technique based on Tauberian theorems for targeting a Lévy-driven system to a pre-specified distribution, where they required the target distribution to be exactly evaluated. Whilst being relevant, the applicability and the impact of these approaches are rather limited in the domain of machine learning.

In this study, we explore the use of Lévy-driven SDEs within MCMC. Encouraged by earlier studies that illustrate the benefits of using heavy-tailed distributions in MCMC (e.g., improved convergence rates) (Stramer & Tweedie, 1999; Jarner & Roberts, 2007), we aim at investigating the potential benefits of the usage of the Lévy motions in LMC, in lieu of the classical Brownian motion. We extend classical LMC and develop a novel Fractional LMC framework, which targets the correct distribution even in the presence of the jumps induced by the Lévy motion. We then develop novel computational methods that can scale up to large-scale problems and provide formal theoretical analysis of the convergence behavior of the proposed approach.

We support our theoretical results by several synthetic and real experiments. Our experiments show that the proposed approach forms a viable alternative to classical LMC with additional benefits, such as providing superior performance in multi-modal settings, higher convergence rates, and robustness to algorithm parameters. The proposed approach also opens up several interesting future directions, as we will point out in Section 5.

2. Technical Background

Stable distributions: Stable distributions are heavy-tailed distributions. They are the limiting distributions in the generalized central limit theorem: the properly scaled sum of i.i.d. random variables, which do *not* necessarily have finite variance, will converge to an α -stable random variable (Samorodnitsky & Taqqu, 1994). In this study, we are interested in the centered *symmetric* α -stable ($\mathcal{S}\alpha\mathcal{S}$) distribution, which is a special case of α -stable distributions.

The $\mathcal{S}\alpha\mathcal{S}$ distribution can be seen as a heavy-tailed generalization of the centered Gaussian distribution. The probability density function (pdf) of an $\mathcal{S}\alpha\mathcal{S}$ distribution cannot be written in closed-form except for certain special cases; however, the characteristic function of the distribution can be written as follows: $X \sim \mathcal{S}\alpha\mathcal{S}(\sigma) \iff \mathbb{E}[\exp(i\omega X)] = \exp(-|\sigma\omega|^\alpha)$. Here, $\alpha \in (0, 2]$ is called the *characteristic exponent* and determines the tail thickness of the distribution: as α gets smaller, $\mathcal{S}\alpha\mathcal{S}$ becomes heavier-tailed. The parameter $\sigma \in \mathbb{R}_+$ is called the *scale* parameter and measures the spread of X around 0.

As an important special case of $\mathcal{S}\alpha\mathcal{S}$, we obtain the Gaussian distribution $\mathcal{S}\alpha\mathcal{S}(\sigma) = \mathcal{N}(0, 2\sigma^2)$ for $\alpha = 2$. In Figure 1, we illustrate the (approximately computed) pdf of the symmetric α -stable distribution for different values of α . As can be clearly observed from the figure, the tails of the distribution vanish quickly when $\alpha = 2$ (i.e. Gaussian), whereas the tails get thicker as we decrease α .

An important property of the α -stable distributions is that their moments can only be defined up to the order α , i.e. $\mathbb{E}[|X|^p] < \infty$ if and only if $p < \alpha$ for $\alpha \in (0, 2)$; implying that X has infinite variance for $\alpha \neq 2$. Moreover, even though the pdf of $\mathcal{S}\alpha\mathcal{S}$ does not admit an analytical form, it is straightforward to draw random samples from stable distributions (Chambers et al., 1976), where efficient implementations are readily available in public software libraries such as the GNU Scientific Library (gnu.org/software/gsl).

SDEs driven by symmetric stable Lévy processes : In this study, we are interested in SDEs driven by symmetric α -stable Lévy processes, which are defined as follows:

$$dX_t = b(X_{t-}, \alpha)dt + dL_t^\alpha, \quad (3)$$

where b is called the drift and is chosen as a function of α in our context, and X_{t-} will be defined in the sequel. Here, L_t^α denotes the D -dimensional α -stable Lévy motion with *independent components*, i.e. each component of L_t^α forms an independent *scalar* α -stable Lévy motion, which is defined as follows for $\alpha \in (0, 2]$ (Duan, 2015):

- (i) $L_0^\alpha = 0$ almost surely.
- (ii) For $t_0 < t_1 < \dots < t_N$, the increments $(L_{t_n}^\alpha - L_{t_{n-1}}^\alpha)$ are independent ($n = 1, \dots, N$).
- (iii) The difference $(L_t^\alpha - L_s^\alpha)$ and L_{t-s}^α have the same distribution: $S_\alpha \mathcal{S}((t-s)^{1/\alpha})$ for $s < t$.
- (iv) L_t^α has *stochastically continuous* sample paths (i.e. continuous in probability): for all $\delta > 0$ and $s \geq 0$, $p(|L_t^\alpha - L_s^\alpha| > \delta) \rightarrow 0$ as $t \rightarrow s$.

Due to the stochastic continuity property, α -stable Lévy motions can have a countable number of discontinuities, which are often referred to as ‘jumps’. As illustrated in Figure 1 (bottom), the size of these jumps becomes larger as α get smaller, since $S_\alpha \mathcal{S}$ becomes heavier tailed. As a consequence, the sample paths of these processes are continuous from the right and they have left limits at every time (Duan, 2015): X_{t-} hence denotes the left limit of X_t at time t . Therefore, these processes are called càdlàg, i.e. the French acronym for ‘*continue à droite, limite à gauche*’.

Similarly to the symmetric α -stable distributions, the symmetric α -stable Lévy motions L_t^α coincide with a scaled Brownian motion $\sqrt{2}B_t$ when $\alpha = 2$. This can be simply verified by observing that the difference $L_t^\alpha - L_s^\alpha$ follows a Gaussian distribution $\mathcal{N}(0, 2(t-s))$ and L_t^α becomes almost surely continuous everywhere.

Riesz potentials and fractional differentiation: Fractional calculus aims to generalize differentiation (and integration) to fractional orders (Herrmann, 2014). The canonical example of fractional differentiation can be given as the *half-derivative* operator, which coincides with the first-order derivative when applied twice to any function.

In this study, we are interested in *fractional Riesz derivatives* (Riesz, 1949), which are closely related to α -stable distributions. The fractional Riesz derivative directly generalizes the *second-order* differentiation to fractional orders and it is a *non-local* operator. In the one dimensional case, it is defined by the following identity:

$$\mathcal{D}^\gamma f(x) \triangleq \mathcal{F}^{-1}\{|\omega|^\gamma \hat{f}(\omega)\}, \quad (4)$$

where \mathcal{F} denotes the Fourier transform and $\hat{f}(\omega) = \mathcal{F}\{f(x)\}$. Here, $\gamma > -1$ is the order of the differentiation: for $\gamma \in (-1, 0)$ we obtain the Riesz potentials¹, which will be our main source of interest, and for $\gamma = 2$ we obtain the

¹For $\gamma < 0$, \mathcal{D}^γ corresponds to fractional integration. However, we follow the fractional calculus literature and still refer to it as fractional differentiation.

usual second-order differentiation up to a sign difference, i.e. $\mathcal{D}^2 f(x) = -d^2 f(x)/dx^2$. Note that $\mathcal{D}^1 f(x)$ does not coincide with first-order differentiation in general.

3. Fractional Langevin Monte Carlo

In this section, we present our main results and construct the proposed Fractional LMC framework step by step. We first develop a Lévy-driven SDE that targets the correct distribution and analyze the weak convergence properties of its Euler discretization. Afterwards, we develop numerical methods for approximate simulation of the proposed SDE and present formal analysis of the approximation error of the numerical schemes and the weak error analysis of the corresponding Euler discretizations.

In the rest of this paper, we restrict α to be in $(1, 2]$ in order the mean of the process to exist. Besides, in all our analyses we focus on the scalar case ($D = 1$) for simplicity; however, all our results can be extended for $D > 1$. All the proofs are provided in the supplementary document.

3.1. Invariant measure and weak convergence analysis

Our first goal is to find a drift b in such a way that the Markov process $(X_t)_{t \geq 0}$ that is a càdlàg solution of the SDE in (3) would have the target distribution π as an invariant distribution. In the following theorem, we present our first main result.

Theorem 1. *Consider the SDE (3), where b is defined as:*

$$b(x, \alpha) \triangleq (\mathcal{D}^{\alpha-2} f_\pi(x)) / \phi(x). \quad (5)$$

Here, $f_\pi(x) \triangleq -\phi(x) \partial_x U(x)$ and \mathcal{D}^γ is defined in (4). Then, π is an invariant measure of the Markov process $(X_t)_{t \geq 0}$ that is a càdlàg solution of the SDE given in (3). Furthermore, if b is Lipschitz continuous, then π is the unique invariant measure of the process $(X_t)_{t \geq 0}$.

The Lipschitz continuity of $\partial_x U$ is a standard condition in LMC for ensuring the uniqueness of the invariant measure, albeit it is often violated in practical applications. In our context, we need b to be Lipschitz continuous for uniqueness, a condition which cannot be easily verified for $\alpha \neq 2$. Here, it is also worth noting that when $\alpha \rightarrow 2$, we obtain the classical Langevin diffusion (1), as $\lim_{\alpha \rightarrow 2} L_t^\alpha = \sqrt{2}B_t$ and $\lim_{\alpha \rightarrow 2} b(x, \alpha) = -\partial_x U(x)$.

Theorem 1 suggests that if we could generate continuous sample paths from $(X_t)_{t \geq 0}$, then we could use them as samples drawn from π . However, this is not possible since the drift b does not admit an analytical form in general, and even if it could be computed exactly, we still could not simulate the SDE (3) exactly as it is a continuous-time process.

For now, let us assume that we can exactly compute the drift b and focus on simulating the SDE by considering

its Euler-Maruyama discretization (Duan, 2015; Panloup, 2008), which is given as follows:

$$\bar{X}_{n+1} = \bar{X}_n + \eta_{n+1} b(\bar{X}_n, \alpha) + \eta_{n+1}^{1/\alpha} \Delta L_{n+1}^\alpha \quad (6)$$

where $n = 1, \dots, N$ denotes the time-steps, N is the total number of time-steps (i.e. iterations), $(\eta_n)_n$ is a sequence of step-sizes, and $(\Delta L_n^\alpha)_n$ is a sequence of i.i.d. standard symmetric α -stable random variables, i.e. $\Delta L_n^\alpha \sim \mathcal{S}\alpha\mathcal{S}(1)$. We can clearly observe that this discretization schema is a fractional generalization of ULA given in (2), where it coincides with ULA when $\alpha = 2$.

The Euler-Maruyama scheme in (6) lets us approximately compute the expectation of a test function g under the target density π , i.e. $\nu(g) \triangleq \int g(X)\pi(dX)$, by using sample averages, given as: $\bar{\nu}_N(g) \triangleq \frac{1}{H_N} \sum_{n=1}^N \eta_n g(\bar{X}_n)$, where $H_N = \sum_{n=1}^N \eta_n$. Even though the convergence properties of the estimators obtained via ULA have been well-established (Roberts & Stramer, 2002; Durmus & Moulines, 2015), it is not clear whether the estimator $\bar{\nu}_N(g)$ converges to the true expectation $\nu(g)$ for $\alpha \neq 2$.

For the convergence analysis, we make use of relatively recent results from the applied probability literature (Panloup, 2008). In order to establish the convergence of our estimator, we need certain conditions to be satisfied. First, we have a rather standard assumption on the step-sizes:

H1. $\lim_{n \rightarrow \infty} \eta_n = 0$, $\lim_{N \rightarrow \infty} H_N = \infty$,

i.e. the step-sizes are required to be decreasing and their sum is required to diverge. Secondly, we need a more technical Lyapunov condition in order to ensure the stochastic process to be mean-reverting.

H2. Let $V : \mathbb{R} \rightarrow \mathbb{R}_+^*$ be a function in \mathcal{C}^2 , satisfying $\lim_{|x| \rightarrow \infty} V(x) = \infty$, $|\partial_x V| \leq C\sqrt{V}$ for some $C > 0$, and $\partial_x^2 V$ is bounded. There exists $a \in (0, 1]$, $\delta > 0$ and $\beta \in \mathbb{R}$, such that $|b|^2 \leq CV^a$ and $b(\partial_x V) \leq \beta - \delta V^a$, where b is defined in (5).

Under these conditions, we present the following corollary to Theorem 1 and (Panloup, 2008, Theorem 2), where we establish the weak convergence of the Euler-Maruyama scheme defined in (6).

Corollary 1. Assume that b is Lipschitz continuous and the conditions **H1** and **H2** hold. If the test function $g = o(V^{p/2+a-1})$ with $p \in (0, 1/2]$, then the following holds:

$$\lim_{N \rightarrow \infty} \bar{\nu}_N(g) = \nu(g), \quad \text{almost surely.}$$

This corollary shows that under certain regularity and Lyapunov conditions, the Euler-Maruyama scheme in (6) still weakly converges for $\alpha \neq 2$, as long as the drift can be computed exactly. Note that we consider the Lipschitz condition for ensuring the uniqueness of the invariant measure;

however, this is not a crucial assumption as one can show that every weak limit of the sequence $\{\bar{\nu}_N\}_N$ is an invariant probability for the SDE in (3).

3.2. Numerical approximation

Even though Corollary 1 ensures the weak convergence of the Euler scheme, its practical implication is somewhat limited since the Riesz derivatives cannot be computed exactly in general. In this section, we develop and analyze numerical methods for approximately computing the drift b .

In (Ortigueira, 2006), it has been shown that for $\gamma \in (-1, 2)$, the Riesz derivative \mathcal{D}^γ of a function $f(x)$ can be defined as the limit of the *fractional centered difference* operator Δ_h^γ , given as: $\mathcal{D}^\gamma f(x) = \lim_{h \rightarrow 0} \Delta_h^\gamma f(x)$, where

$$\Delta_h^\gamma f(x) \triangleq (1/h^\gamma) \sum_{k=-\infty}^{\infty} g_{\gamma,k} f(x - kh), \quad (7)$$

and $g_{\gamma,k} \triangleq (-1)^k \Gamma(\gamma + 1) / (\Gamma(\frac{\gamma}{2} - k + 1) \Gamma(\frac{\gamma}{2} + k + 1))$. By using the above definition, we can rewrite our drift as: $b(x, \alpha) = (\lim_{h \rightarrow 0} \Delta_h^{\alpha-2} f_\pi(x)) / \phi(x)$, where $f_\pi(x)$ is defined in Theorem 1.

We now propose our first numerical scheme for approximating the drift b by following Çelik & Duman (2012):

$$b(x, \alpha) \approx \tilde{b}_{h,K}(x, \alpha) \triangleq (\Delta_{h,K}^{\alpha-2} f_\pi(x)) / \phi(x), \quad (8)$$

where $\Delta_{h,K}^\gamma$ is the *truncated* fractional central difference operator, defined as follows:

$$\Delta_{h,K}^\gamma f(x) \triangleq (1/h^\gamma) \sum_{k=-K}^K g_{\gamma,k} f(x - kh). \quad (9)$$

Here, we merely replaced the Riesz derivative with the central difference operator where we fixed h and truncated the infinite summation in order the numerical scheme to be computationally tractable. We provide a numerically stable implementation of (8) in the supplementary document.

The scheme in (8) provides us a practical way for approximately computing the drift. However, for fixed h and K , this approach would yield a certain approximation error and therefore Corollary 1 would no longer hold if we replace b by \tilde{b} in (6). Throughout this section, we analyze this approximation error and the weak error of the Euler scheme with the approximate drift.

We first analyze the approximation error of our numerical scheme in (8). Since $\phi(x)$ is constant for a given x , we focus on $|\mathcal{D}^\gamma f_\pi(x) - \Delta_{h,K}^\gamma f_\pi(x)|$. Here, we first need a technical regularity condition on f_π .

H3. $f_\pi(x) \in \mathcal{C}^3(\mathbb{R})$ and all derivatives up to order three belong to $\mathcal{L}_1(\mathbb{R})$.

We need an additional assumption on f_π , which ensures the tails of the target distribution π vanish sufficiently quickly.

H4. $|f_\pi(x - kh)| \leq C \exp(-|k|h)$ for some $C > 0$ and $|k| > K$ for some $K \in \mathbb{N}_+$.

Now, we present our second main result.

Theorem 2. *Assume that the conditions **H3** and **H4** hold. Then, for $\gamma \in (-1, 0)$, the following bound holds:*

$$|\mathcal{D}^\gamma f_\pi(x) - \Delta_{h,K}^\gamma f_\pi(x)| = \mathcal{O}(h^2 + 1/(hK)), \quad (10)$$

as h goes to zero.

Theorem 2 shows that the error induced by our numerical approximation scheme is bounded and can be made arbitrarily small by decreasing h and increasing K . We can also observe that for fixed K , the optimal $h = \mathcal{O}(K^{-1/3})$.

The hidden constant in the right hand side of (10) is allowed to depend on x and let it be denoted as $C(x)$. In order to ease the analysis, in the rest of the paper we will assume that $\sup_x C(x)/\phi(x) < \infty$, so that Theorem 2 can be directly used for bounding the error $|b - \tilde{b}|$ for any x . Note that this is a mild assumption and holds trivially when X belongs to a bounded domain (e.g. the setting in (Wang et al., 2015)).

We now consider the following Euler-Maruyama discretization of (3) with the approximate drift:

$$\tilde{X}_{n+1} = \tilde{X}_n + \eta_{n+1} \tilde{b}_{h,K}(\tilde{X}_n, \alpha) + \eta_{n+1}^{1/\alpha} \Delta L_{n+1}^\alpha, \quad (11)$$

where the corresponding estimator is defined as: $\tilde{\nu}_N(g) \triangleq (1/H_N) \sum_{n=1}^N \eta_n g(\tilde{X}_n)$. Even in this approximate Euler-scheme, we still obtain ULA as a special case of (11), as we have $\lim_{\alpha \rightarrow 2} \tilde{b}_{h,K}(x, \alpha) = -\partial_x U(x)$.

As opposed to $\bar{\nu}_N(g)$, $\tilde{\nu}_N(g)$ does not converge to $\nu(g)$ due to the error induced by the numerical approximation. However, fortunately, the weak error of this Euler scheme can still be bounded, as we show in the following theorem. For this result, we need an additional ergodicity condition².

H5. *The SDE (3) and $dX_t = \tilde{b}_{h,K}(X_t, \alpha)dt + dL_t^\alpha$ are geometrically ergodic with their unique invariant measures.*

Theorem 3. *Assume that the conditions **H1**, **3** and **5** hold, **H2** holds for \tilde{b} , and K is chosen in such a way that **H4** holds for any x . Finally, assume that $|\partial_x g|$ is bounded. Then, the following bound holds almost surely:*

$$|\nu(g) - \lim_{N \rightarrow \infty} \tilde{\nu}_N(g)| = \mathcal{O}(h^2 + 1/(hK)). \quad (12)$$

This theorem shows that the weak error of the discretization in (11) is dominated by the numerical error induced by \tilde{b} and can be made arbitrarily small by tuning h and K .

²Proving the ergodicity of the SDEs in **H5** is beyond the scope of this study; more information can be found in (Masuda, 2007).

3.3. Multidimensional case

Even though we have focused on the scalar case (i.e. $D = 1$) so far, we can generalize the presented results to vector processes by using the same proof strategies since the components of L_t^α are independent³. For $D > 1$, the drift turns out to be a multidimensional generalization of (5) and has the following form: (for $d = 1, \dots, D$)

$$[b(x, \alpha)]_d = \mathcal{D}_{x_d}^{\alpha-2} \{-\phi(x) \partial_{x_d} U(x)\} / \phi(x), \quad (13)$$

where $[v]_i$ denotes the i 'th component of a vector v and $\mathcal{D}_{x_d}^\alpha$ denotes the *partial* fractional Riesz derivative along the direction x_d (Ortigueira et al., 2014). With this definition of the multidimensional drift, similar to the scalar case, we obtain the classical Langevin equation as a special case of (3), since $\lim_{\alpha \rightarrow 2} b(x, \alpha) = -\nabla U(x)$ for $D > 1$.

In applications, we can approximate (13) by applying the same numerical technique presented in (8) to each dimension d . However, for large D , this approach would be impractical since it would require the fractional derivatives to be computed D times at each iteration.

In this section, we propose a second scheme for approximating the fractional Riesz derivatives. The current approach is a computationally more efficient variant of the first numerical scheme presented in (7) and it is given as follows: $\mathcal{D}^\gamma f_\pi(x) \approx g_{\gamma,0} f_\pi(x)$, where $g_{\gamma,0} = \Gamma(\gamma + 1)/\Gamma(\frac{\gamma}{2} + 1)^2$ for $x \in \mathbb{R}$. In other words, we approximate the fractional derivatives by using only the first term of the centered difference operator defined in (7). When all the partial fractional derivatives in (13) are approximated with this approach, the multidimensional drift greatly simplifies and has the following form: (for $D > 1$)

$$b(X, \alpha) \approx \hat{b}(X, \alpha) \triangleq -c_\alpha \nabla U(X), \quad (14)$$

where $c_\alpha \triangleq \Gamma(\alpha - 1)/\Gamma(\alpha/2)^2$. We finally consider a discretization of (3) where the drift is approximated by (14) and ultimately propose the Fractional Langevin Algorithm (FLA), defined as follows:

$$\hat{X}_{n+1} = \hat{X}_n - \eta_{n+1} c_\alpha \nabla U(\hat{X}_n) + \eta_{n+1}^{1/\alpha} \Delta L_{n+1}^\alpha. \quad (15)$$

Similar to the previous discretization schemes given in (6) and (11), FLA generalizes ULA as well, since $\lim_{\alpha \rightarrow 2} \hat{b}(x, \alpha) = -\nabla U(x)$. Besides, FLA has the exact same computational complexity as ULA, since it only requires to compute ∇U and generate ΔL_n^α . Another interesting observation is that c_α increases as α decreases,

³While extending our results to $D > 1$, the independence of the components of L_t^α turns out to be a crucial requirement since the spectral measure of the corresponding multivariate stable distribution becomes discrete (Nolan, 2008). Our results cannot be directly extended to SDEs that are driven by other multivariate stable processes, such as isotropic stable processes (Nolan, 2013).

implying that FLA tends to increase the ‘weight’ of the gradient as the driving process becomes heavier-tailed.

We now present our last theoretical result where we analyze the approximation error of the simplified scheme for $D = 1$, and present it as a corollary to Theorems 2 and 3.

Corollary 2. *Assume that the conditions **H3** and **4** hold. Let $K_x \in \mathbb{N}_+$ be a value that satisfies **H4** for a given x , and let $r : \mathbb{R} \rightarrow \mathbb{R}_+$ be a function defined as:*

$$r(x) \triangleq \left| \sum_{k \in \llbracket -K_x, K_x \rrbracket \setminus 0} \frac{g_{\gamma, k} f_\pi(x - kh)}{g_{\gamma, 0} f_\pi(x)} + 1 \right|^{1/\gamma}.$$

Then, for $-1 < \gamma < 0$, the following bound holds:

$$\left| \mathcal{D}^\gamma f_\pi(x) - g_{\gamma, 0} f_\pi(x) \right| = \mathcal{O}(r^2(x) + 1/(r(x)K_x)).$$

Furthermore, if **H1** holds, **H2** and **5** hold for \hat{b} , and $|\partial_x g|$ is bounded, then, the following bound holds:

$$|\nu(g) - \lim_{N \rightarrow \infty} \hat{\nu}_N(g)| = \mathcal{O}(r^2(x^*) + 1/(r(x^*)K_{x^*}))$$

almost surely, where $x^* = \arg \max_x [r^2(x) + 1/(r(x)K_x)]$ and $\hat{\nu}_N(g) \triangleq (1/H_N) \sum_{n=1}^N \eta_n g(\hat{X}_n)$.

Here, the term $r(x)$ plays a similar role as the parameter h in (7). This corollary shows that the approximation quality of (14) may vary depending on the particular x where $\mathcal{D}^\gamma f_\pi(x)$ is evaluated, and depending on the values of K_x and $r(x)$, \hat{b} might even provide more accurate approximations than \tilde{b} does. As a result, we observe that the weak error is dominated by the largest numerical error induced by \hat{b} . On the other hand, even if \hat{b} would have a higher approximation error when compared to \tilde{b} , we would expect that the scheme in (15) to be better behaved than (11), since it is less prone to numerical instability.

3.4. Large-scale Bayesian posterior sampling

In Bayesian machine learning, the target distribution π is often chosen as the Bayesian posterior: $\pi(X) = p(X|Y)$, where $Y \equiv \{Y_i\}_{i=1}^{N_Y}$ is a set of observed i.i.d. data points. This choice of the target distribution imposes the following form on the potential energy: $U(X) = -(\sum_{i=1}^{N_Y} \log p(Y_i|X) + \log p(X))$, where $p(Y_i|X)$ is the likelihood function and $p(X)$ is the prior distribution.

In large scale applications, N_Y becomes very large and therefore computing ∇U at each iteration can be computationally prohibitive. Inspired by the Stochastic Gradient Langevin Dynamics (SGLD) algorithm (Welling & Teh, 2011), which extends ULA to large-scale settings, we extend FLA by replacing the exact gradients ∇U in (15) with an unbiased estimator, given as follows:

$$\nabla \tilde{U}_n(X) = -[\nabla \log p(X) + \frac{N_Y}{N_\Omega} \sum_{i \in \Omega_n} \nabla \log p(Y_i|X)],$$




Figure 2. Top: the double-well potential. Middle: the empirical distribution obtained via ULA (corresponds to FLA with $\alpha = 2$). Bottom: the empirical distribution obtained via FLA ($\alpha = 1.7$).

where $\Omega_n \subset \{1, \dots, N_Y\}$ is a random data subsample that is drawn with replacement at iteration n and $N_\Omega \ll N_Y$ denotes the number of elements in Ω . We call the resulting algorithm Stochastic Gradient FLA (SG-FLA). Note that SG-FLA coincides with SGLD when $\alpha = 2$. We leave the convergence analysis of SG-FLA as a future work.

4. Experiments

The double-well potential: We conduct our first set of experiments on a synthetic setting where we consider the double-well potential, defined as follows:

$$U(x) = (x + 5)(x + 1)(x - 1.02)(x - 5)/10 + 0.5.$$

We illustrate the double-well potential in Figure 2 (top). It can be observed that the potential contains two well-separated modes with different heights, which makes the problem challenging.

In our first experiment, we consider our first discretization scheme presented in (11), where we approximate the true drift b , by \hat{b} . Here, we use decreasing step-sizes that are determined as $\eta_n = (a_\eta/n)^{b_\eta}$, where we fix $a_\eta = 10^{-7}$ and $b_\eta = 0.6$. In each experiment, we generate $N = 5000$ samples by using (11) and estimate the mean of the target distribution π by using the sample average. For each α , we run this scheme for different values of h and K , repeat each experiment 5 times, and monitor the bias, where the ground truth is obtained via a numerical integrator.

We first fix $h = 0.06$ and monitor the bias for increasing values of K . Here, we define the notion of an optimal K as the smallest K , for whose larger values the performance improvement becomes negligible. As we can observe from Figure 3 (top), the bias is gracefully degrading for increasing K , where the optimal K depends on the choice of α . We also observe that, modest values of K seem sufficient for obtaining accurate results, especially for small α .




Figure 3. Evaluation of the scheme in (11) on the estimation of the mean of the double-well potential. Top: the average bias vs K for a fixed h . Bottom: the average bias vs h for a fixed K .

In our second experiment, we fix $K = 15$ and monitor the bias for different values of h . The results are illustrated in Figure 3 (bottom). We observe that the results support our theory: for very small values of h , the term $1/(hK)$ in the bound of Theorem 3 dominates since K is fixed. Therefore, we observe a drop in the bias as we increase h up to a certain point, and then the bias gradually increases along with h . The results show that the performance becomes more sensitive to the value of h , as α becomes smaller.

Even though the results in Figure 3 are promising, in practical applications we would not be able to use the scheme in (8) due to computational issues. In our next experiment, we aim to assess the approximation error of our second approximation scheme \hat{b} given in (14). However, the error $|b(x, \alpha) - \hat{b}(x, \alpha)|$ cannot be measured in a straightforward manner, since b cannot be computed exactly and the error itself depends on the particular point x .

Here, we develop an intuitive accuracy criterion for getting better insight into this error, where the aim is to compute the value of K for which $\tilde{b}_{h,K}(x, \alpha)$ and $\hat{b}(x, \alpha)$ would yield similar approximation errors on average. For a given x and fixed h , we first choose a large enough $K^* \in \mathbb{N}_+$ and compute $b^*(x, \alpha) \triangleq \tilde{b}_{h,K^*}(x, \alpha)$ as our reference for $b(x, \alpha)$. Then, for $K \in \llbracket 1, K^* \rrbracket$, we compute the approximation error $e(x, \alpha, K) \triangleq |\tilde{b}_{h,K}(x, \alpha) - b^*(x, \alpha)|$ and the error induced by the ultimate approximation scheme: $\hat{e}(x, \alpha) \triangleq |\hat{b}(x, \alpha) - b^*(x, \alpha)|$. We then find the value of K for which $e(x, \alpha, K)$ and $\hat{e}(x, \alpha)$ are the closest: $\kappa(x, \alpha) = \arg \min_K |e(x, \alpha, K) - \hat{e}(x, \alpha)|$. We finally evaluate κ on I different points $\{x_i\}_{i=1}^I$ and use the average of these values as the measure of accuracy of \hat{b} , defined as: $\hat{\kappa}(\alpha) = (1/I) \sum_{i=1}^I \kappa(x_i, \alpha)$. Intuitively, this value is expected to be large when \hat{b} yields a low error.

In order to assess the accuracy of \hat{b} in the double-well problem, we compute $\hat{\kappa}(\alpha)$ for different values of α , where we fix $h = 0.06$, $K^* = 170$, and choose $\{x_i\}_i$ as $I = 200$ evenly-spaced points from the interval $[-5, 5]$. The results are given in Table 4. The results show that, despite its sim-

Table 1. Evaluation of the accuracy of \hat{b} .

	$\alpha = 1.5$	$\alpha = 1.6$	$\alpha = 1.7$	$\alpha = 1.8$	$\alpha = 1.9$
$\hat{\kappa}(\alpha)$	19.31	14.12	12.72	8.64	7.03




Figure 4. The average bias obtained via FLA for different values of α . Note that FLA corresponds to ULA when $\alpha = 2$.

plicity, \hat{b} is able to provide reasonably accurate estimates for b . We observe that for $\alpha = 1.6$, $\hat{\kappa}(\alpha)$ becomes 14.12, which is even larger than the optimal K for $\alpha = 1.6$, as shown in Figure 3. Therefore, it is promising to use \hat{b} in real applications since it yields sufficiently accurate approximations with less computational requirements and does not require additional tuning for h and K .

In our last experiment on the double-well potential, we evaluate the ultimately proposed approach FLA on estimation of the mean of the target distribution. Similarly to the previous experiments, we run FLA for different values of α , where we try several values for the hyper-parameters a_η and b_η and report the best results for each α . In each experiment, we generate $N = 50000$ samples by using (15) and repeat the procedure 10 times. We first illustrate two typical empirical distributions obtained via FLA and ULA in Figure 2 (middle, bottom). It can be clearly observed that ULA can locate only one of the modes, whereas FLA is able to locate both of the modes, thanks to the jumps of the α -stable processes. This circumstance also reflects in the average bias, as illustrated in Figure 4. The results show that for $\alpha = 2$ the average bias is around 3, implying that the algorithm concentrates on either one of the modes at each trial, whereas we observe that the bias rapidly decreases as we decrease α . The best performance is achieved when $\alpha = 1.75$. Finally we note that these results are also in line with the best-performing results given in Figure 3.

Matrix factorization: In our second set of experiments, we switch to a large-scale Bayesian machine learning context. We explore the use of SG-FLA on a large-scale link prediction application where we consider the following probabilistic matrix factorization model (Gemulla et al., 2011; Salakhutdinov & Mnih, 2008): $A_{il} \sim \mathcal{N}(0, 1)$, $B_{lj} \sim \mathcal{N}(0, 1)$, $Y_{ij}|A, B \sim \mathcal{N}(\sum_l A_{il}B_{lj}, 1)$, where $Y \in \mathbb{R}^{I \times L}$ is the observed data matrix with possible missing entries, and $A \in \mathbb{R}^{I \times L}$ and $B \in \mathbb{R}^{L \times J}$ are the latent factor matrices. The aim in this application is to predict the missing values of Y by using a low-rank approximation. Recently, SGLD has been proven successful on similar models (Ahn et al., 2015; Şimşekli et al., 2015;

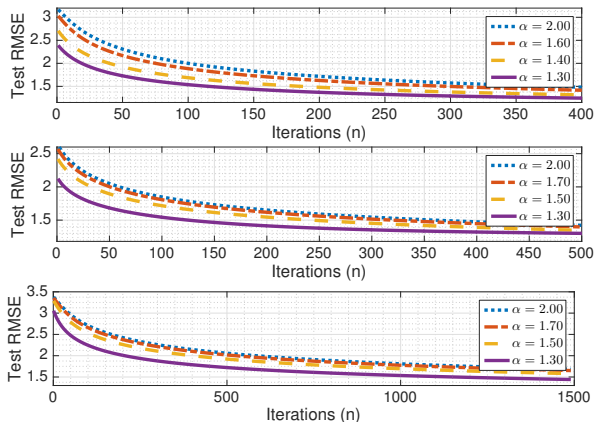


Figure 5. The performance of SG-FLA on a link prediction application. Top: ML-1M, middle: ML-10M, bottom: ML-20M.

Durmus et al., 2016; Şimşekli et al., 2017).

In this set of experiments, we apply SG-FLA on the three MovieLens movie ratings datasets (grouplens.org): MovieLens 1Million (ML-1M), 10Million (ML-10M), and 20Million (ML-20M). The ML-1M dataset contains 1 million non-zero entries, where $I = 3883$ (movies) and $J = 6040$ (users). The ML-10M dataset contains 10 million non-zero entries, where $I = 10681$ and $J = 71567$. Finally, the ML-20M dataset contains 20 million ratings, where $I = 27278$ and $J = 138493$. In our experiments, we randomly select 10% of the data as the test set and use the remaining data for generating the samples. The rank L is set to 10 for all datasets. In all experiments, we use decreasing step-sizes, where we fix $b_\eta = 0.51$ and try several values for a_η and report the best results. We set $N_\Omega = N_Y/10$ where N_Y denotes the number of non-zero entries in a given dataset.

Figure 5 shows the root mean squared-errors (RMSE) that are obtained on the three test sets. In all these experiments, we observe that the rate of convergence of SG-FLA increases as we decrease α from 2.0 (i.e. SGLD) to 1.3. In the case when $\alpha < 1.3$, the jumps induced by the stable-Lévy motion becomes very large and the performance starts degrading. These results show that SG-FLA can be considered as a viable alternative to SGLD in large scale settings and it can provide improved performance over SGLD via minor algorithmic modifications, which come with the expense of tuning an additional parameter α .

Sigmoid Belief Networks: In our last set of experiments, we investigate the use of SG-FLA on Sigmoid Belief Networks (SBN) (Gan et al., 2015), which have been investigated in recent Stochastic Gradient MCMC studies (Chen et al., 2015). We make use of the software provided in (Chen et al., 2015) and employ the identical experimental setup described therein: the binary observed data are assumed to be generated from a single binary hidden layer with sigmoid activations. The overall model is applied on




Figure 6. The test log-likelihoods obtained via SG-FLA on SBNs, as a function of step-size η .

the MNIST dataset, which contains 70K binary images (of size 28×28) corresponding to different digits.

In our experiments, we use an SBN with 100 hidden units. We use a training set of 60K images and 10K images for testing, set the size of the data subsample $N_\Omega = 200$, and run SG-FLA for 5000 iterations for training. Finally, we estimate the test likelihoods by using an annealed importance sampler (Salakhutdinov & Murray, 2008).

As opposed to our previous experiments, we use constant step-sizes in these experiments, i.e. $\eta_n = \eta$ for all n , and investigate the performance of SG-FLA on SBNs for different values of η and α . The results are illustrated in Figure 6. We can observe that for small values of η , SG-FLA yields similar test likelihoods for all values of α . However, as we increase the step size, we observe that the test likelihood of SGLD ($\alpha = 2$) starts to diverge, whereas SG-FLA becomes more robust to large step sizes as α gets smaller. When $\alpha = 1.6$ the test likelihood stays almost constant for increasing values of η . We do not observe an improvement in the performance for $\alpha < 1.6$.

5. Conclusion and Future Directions

In this study, we explored the use of Lévy-driven SDEs within MCMC and presented a novel FLMC framework. We first showed that FLMC targets the correct distribution and then developed novel and scalable computational methods for practical applications. We provided formal analysis of the convergence properties and the approximation quality of the proposed numerical schemes. We supported our theory with several experiments, which showed that FLMC brings various benefits, such as providing superior performance in multi-modal settings, higher convergence rates, and robustness to algorithm parameters.

The proposed framework opens up several interesting future directions: (i) the use of FLMC in simulated annealing for global optimization (Chen et al., 2016), where the jumps might bring further advantages (ii) extension of FLMC to ‘stable-like’ processes (Bass, 1988), where α can depend on X (iii) incorporation of the local geometry for faster convergence (Patterson & Teh, 2013; Li et al., 2016; Şimşekli et al., 2016a) (iv) the use of SG-FLA in Bayesian model selection (Şimşekli et al., 2016b).

Acknowledgments

The author would like to thank to Alain Durmus for his helps on the proofs, and to Roland Badeau, A. Taylan Cemgil, and Gaël Richard for fruitful discussions. The author would also like to thank to Changyou Chen for sharing the code used in the experiments conducted on SBNs. This work is partly supported by the French National Research Agency (ANR) as a part of the FBIMATRIX project (ANR-16-CE23-0014), and the EDISON 3D project (ANR-13-CORD-0008-02).

References

- Ahn, S., Korattikara, A., Liu, N., Rajan, S., and Welling, M. Large-scale distributed Bayesian matrix factorization using stochastic gradient MCMC. In *KDD*, 2015.
- Bass, R. F. Uniqueness in law for pure jump Markov processes. *Probability Theory and Related Fields*, 79(2): 271–287, 1988.
- Çelik, C. and Duman, M. Crank–Nicolson method for the fractional diffusion equation with the Riesz fractional derivative. *Journal of Computational Physics*, 231(4): 1743–1750, 2012.
- Chambers, J. M., Mallows, C. L., and Stuck, B. W. A method for simulating stable random variables. *Journal of the american statistical association*, 71(354):340–344, 1976.
- Chen, C., Ding, N., and Carin, L. On the convergence of stochastic gradient MCMC algorithms with high-order integrators. In *Advances in Neural Information Processing Systems*, pp. 2269–2277, 2015.
- Chen, C., Carlson, D., Gan, Z., Li, C., and Carin, L. Bridging the gap between stochastic gradient MCMC and stochastic optimization. In *AISTATS*, 2016.
- Şimşekli, U., Badeau, R., Cemgil, A. T., and Richard, G. Stochastic quasi-Newton Langevin Monte Carlo. In *ICML*, 2016a.
- Şimşekli, U., Badeau, R., Richard, G., and Cemgil, A. T. Stochastic thermodynamic integration: efficient Bayesian model selection via stochastic gradient MCMC. In *ICASSP*, 2016b.
- Şimşekli, U., Durmus, A., Badeau, R., Richard, G., Moulines, E., and Cemgil, A. T. Parallelized stochastic gradient Markov Chain Monte Carlo algorithms for non-negative matrix factorization. In *ICASSP*, 2017.
- Ditlevsen, P. D. Anomalous jumping in a double-well potential. *Physical Review E*, 60(1):172, 1999.
- Duan, J. *An Introduction to Stochastic Dynamics*. Cambridge University Press, New York, 2015.
- Durmus, A. and Moulines, E. Non-asymptotic convergence analysis for the unadjusted Langevin algorithm. *arXiv preprint arXiv:1507.05021*, 2015.
- Durmus, A., Şimşekli, U., Moulines, E., Badeau, R., and Richard, G. Stochastic gradient Richardson-Romberg Markov Chain Monte Carlo. In *NIPS*, 2016.
- Eliazar, I. and Klafter, J. Lévy-driven Langevin systems: Targeted stochasticity. *Journal of statistical physics*, 111(3):739–768, 2003.
- Gan, Z., Henao, R., Carlson, D. E., and Carin, L. Learning deep sigmoid belief networks with data augmentation. In *AISTATS*, volume 38, pp. 268–276, 2015.
- Gemulla, R., Nijkamp, E., J., Haas, P., and Sismanis, Y. Large-scale matrix factorization with distributed stochastic gradient descent. In *ACM SIGKDD*, 2011.
- Herrmann, R. *Fractional calculus: an introduction for physicists*. World Scientific, 2014.
- Jarner, S. F. and Roberts, G. O. Convergence of heavy-tailed Monte Carlo Markov Chain algorithms. *Scandinavian Journal of Statistics*, 34(4):781–815, 2007.
- Konakov, V. and Menozzi, S. Weak error for stable driven stochastic differential equations: Expansion of the densities. *Journal of Theoretical Probability*, 24(2):454–478, 2011.
- Kuruoglu, E. E. *Signal processing in α -stable noise environments: a least l_p -norm approach*. PhD thesis, University of Cambridge, 1999.
- Lamberton, D. and Pages, G. Recursive computation of the invariant distribution of a diffusion: the case of a weakly mean reverting drift. *Stochastics and dynamics*, 3(04): 435–451, 2003.
- Lévy, P. *Théorie de l’addition des variables aléatoires*. Gauthiers-Villars, Paris, 1937.
- Li, C., Chen, C., Carlson, D., and Carin, L. Preconditioned stochastic gradient Langevin dynamics for deep neural networks. In *AAAI Conference on Artificial Intelligence*, 2016.
- Mandelbrot, B. B. *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk. Selecta Volume E*. Springer Science & Business Media, 2013.
- Masuda, H. Ergodicity and exponential β -mixing bounds for multidimensional diffusions with jumps. *Stochastic processes and their applications*, 117(1):35–56, 2007.

- Mikulevičius, R. and Zhang, C. On the rate of convergence of weak Euler approximation for nondegenerate SDEs driven by Lévy processes. *Stochastic Processes and their Applications*, 121(8):1720–1748, 2011.
- Nolan, J. P. An overview of multivariate stable distributions. *Technical Report*, 2008.
- Nolan, J. P. Multivariate elliptically contoured stable distributions: theory and estimation. *Computational Statistics*, 28(5):2067–2089, 2013.
- Ortigueira, M. D. Riesz potential operators and inverses via fractional centred derivatives. *International Journal of Mathematics and Mathematical Sciences*, 2006, 2006.
- Ortigueira, M. D., Laleg-Kirati, T. M., and Machado, J. A. T. Riesz potential versus fractional Laplacian. *Journal of Statistical Mechanics*, (09), 2014.
- Panloup, F. Recursive computation of the invariant measure of a stochastic differential equation driven by a Lévy process. *The Annals of Applied Probability*, 18(2):379–426, 2008.
- Patterson, S. and Teh, Y. W. Stochastic gradient Riemannian Langevin dynamics on the probability simplex. In *Advances in Neural Information Processing Systems*, 2013.
- Riesz, M. L'intégrale de Riemann-Liouville et le problème de Cauchy. *Acta mathematica*, 81(1):1–222, 1949.
- Roberts, G. O. and Stramer, O. Langevin Diffusions and Metropolis-Hastings Algorithms. *Methodology and Computing in Applied Probability*, 4(4):337–357, December 2002. ISSN 13875841.
- Rosky, P. J., Doll, J. D., and Friedman, H. L. Brownian dynamics as smart Monte Carlo simulation. *The Journal of Chemical Physics*, 69(10):4628–4633, 1978.
- Salakhutdinov, R. and Mnih, A. Bayesian probabilistic matrix factorization using Markov Chain Monte Carlo. In *ICML*, pp. 880–887, 2008.
- Salakhutdinov, R. and Murray, I. On the quantitative analysis of deep belief networks. In *ICML*, pp. 872–879, 2008.
- Samorodnitsky, G. and Taqqu, M. S. *Stable non-Gaussian random processes: stochastic models with infinite variance*, volume 1. CRC press, 1994.
- Şimşekli, U., Koptagel, H., Gültaş, H., Cemgil, A. T., Öztoprak, F., and Birbil, Ş. İ. Parallel stochastic gradient Markov Chain Monte Carlo for matrix factorisation models. *arXiv preprint arXiv:1506.01418*, 2015.
- Stramer, O. and Tweedie, R. L. Langevin-type models ii: self-targeting candidates for MCMC algorithms. *Methodology and Computing in Applied Probability*, 1(3):307–328, 1999.
- Wang, Y. X., Fienberg, S. E., and Smola, A. J. Privacy for free: Posterior sampling and stochastic gradient Monte Carlo. In *ICML*, pp. 2493–2502, 2015.
- Welling, M. and Teh, Y. W. Bayesian learning via stochastic gradient Langevin dynamics. In *International Conference on Machine Learning*, pp. 681–688, 2011.
- Yanovsky, V. V., Chechkin, A. V., Schertzer, D., and Tur, A. V. Lévy anomalous diffusion and fractional Fokker-Planck equation. *Physica A: Statistical Mechanics and its Applications*, 282(1):13–34, 2000.