Federated Sampling with Langevin Algorithm under Isoperimetry

Anonymous authors
Paper under double-blind review

Abstract

Federated learning uses a set of techniques to efficiently distribute the training of a machine learning algorithm across several devices, who own the training data. These techniques critically rely on reducing the communication cost—the main bottleneck—between the devices and a central server. Federated learning algorithms usually take an optimization approach: they are algorithms for minimizing the training loss subject to communication (and other) constraints. In this work, we instead take a Bayesian approach for the training task, and propose a communication-efficient variant of the Langevin algorithm to sample *a posteriori*. The latter approach is more robust and provides more knowledge of the *a posteriori* distribution than its optimization counterpart. We analyze our algorithm without assuming that the target distribution is strongly log-concave. Instead, we assume the weaker log Sobolev inequality, which allows for nonconvexity.

1 Introduction

Federated learning uses a set of techniques to efficiently distribute the training of a machine learning algorithm across several devices, who own the training data. In a federated learning environment, several considerations are considered, such as privacy or communication between the devices (Kairouz et al., 2021). In particular, the communication cost between the devices and a central server (Konečnỳ et al., 2016; McMahan et al., 2017) is one of the main bottlenecks in a federated learning environment. Indeed, classical training algorithms such as the Stochastic Gradient Descent (SGD) can formally be distributed between the server and the devices. But this requires communications between the devices and the server, raising communication concerns. Therefore, most federated learning algorithms rely on compressing the messages that are exchanged between the server and the devices. For instance, variants SGD involving compression (Karimireddy et al., 2020; Horváth et al., 2019b; Haddadpour et al., 2021; Das et al., 2020; Gorbunov et al., 2021) have successfully been applied in several engineering setups (Kairouz et al., 2021; Dimitriadis et al., 2022).

Sampling in a federated learning environment. Most of the literature on the training of federated learning learning algorithm focused on the optimization paradigm, i.e., the training algorithm is actually a communication efficient optimization algorithm to minimize the training loss F. But, optimization is not the only paradigm in the training of learning algorithms. The Bayesian approach relies on sampling from a distribution π proportional to $\exp(-F)$ rather than minimizing F. In particular, the Bayesian approach provides a number of samples from $\exp(-F)$ rather than a single minimizer of F. The distribution π is often the the posterior distribution of some Bayesian model. In this case, sampling from π allows to perform some Bayesian computations such as computing confidence intervals.

In this paper, we take a Bayesian approach to federated learning, and similarly to Vono et al. (2022); Deng et al. (2021); El Mekkaoui et al. (2021), we focus on the communication complexity (we ignore privacy concerns). We introduce a communication efficient variant of the Langevin algorithm, a widely used algorithm to sample from a target distribution $\pi \propto \exp(-F)$ where F is a smooth nonconvex function but π satisfies the log-Sobolev inequality (Villani, 2009). We study the complexity of our communication efficient algorithm to sample from π by exploiting the connections between optimization and sampling. More precisely, our approach has roots in the recent literature on viewing Langevin algorithm as an optimization algorithm

over a space of probability distributions called the Wasserstein space (Wibisono, 2018; Durmus et al., 2019; Ambrosio et al., 2005).

1.1 Related Work

Most approaches to train federated learning algorithms rely on minimizing the training loss with a communication efficient variant of SGD. Several training algorithms have been proposed, but we mention Karimireddy et al. (2020); Horváth et al. (2019b); Haddadpour et al. (2021); Das et al. (2020); Gorbunov et al. (2021) because these papers contain the state of the art results in terms of minimization of the training loss. In this work we shall specifically use the MARINA gradient estimator introduced in Gorbunov et al. (2021) and inspired from Nguyen et al. (2017).

The literature on sampling with Langevin algorithm is also large. In recent years, the machine learning research community has specifically been interested in the complexity of Langevin algorithm (Dalalyan, 2017; Durmus & Moulines, 2017). In the case where F is not convex, recent works include Vempala & Wibisono (2019); Wibisono (2019); Ma et al. (2019); Chewi et al. (2021) which use an isoperimetric-type inequality (Villani, 2009, Chapter 21) such as the log Sobolev inequality to prove convergence in Kullback-Leibler divergence. Other works on Langevin in the nonconvex case include Balasubramanian et al. (2022); Cheng et al. (2018); Majka et al. (2020); Mattingly et al. (2002).

The closest papers to our work are Vono et al. (2022); Deng et al. (2021). Like our work, these papers study the convergence of an efficient variant of Langevin algorithm for federated learning. The key difference between these papers and our work is that they assume F strongly convex, whereas we only assume that the target distribution $\pi \propto \exp(-F)$ satisfies the log Sobolev inequality (LSI), which allows F to be nonconvex. More precisely, the strong convexity of F implies LSI (Villani, 2009, Chapter 21), but LSI allows F to be nonconvex. We compare our complexity results to their in Table 1.

Paper	Algorithm	Communication	Assumption
Vono et al. (2022)	QLSD^{++}	**Reduced**	Strong convexity
Deng et al. (2021)	(FA-LD)	**Reduced**	Strong convexity
Wibisono (2019)	Langevin	Expensive	**log Sobolev inequality**
This paper	Langevin-Marina	**Reduced**	**log Sobolev inequality**

Table 1: Langevin-Marina is the only algorithm in the list that works under LSI only, with a reduced communication complexity. The four algorithms in the list have the same complexity $\tilde{\mathcal{O}}\left(\frac{d}{\varepsilon^2}\right)$ to achieve ε accuracy in 2-Wasserstein distance in dimension d.

Our contributions are summarized below.

1.2 Contributions

We consider the problem of sampling from $\pi \propto \exp(-F)$, where π satisfies LSI, in a federated learning setup. We make the following contributions:

- We propose a communication efficient variant of Langevin algorithm, called Langevin-MARINA, that can be distributively implemented between the central server and the devices. Langevin-MARINA

relies on compressed communications. More precisely, we borrow the MARINA estimator used in optimization to perform Sampling.

- We analyze the complexity of our sampling algorithm in terms of the Kullback-Leibler divergence, the Total Variation distance and the 2-Wasserstein distance. Our approach relies on viewing our sampling problem as an optimization problem over a space of probability measures, and allows F to be nonconvex.
- Our sampling algorithm, Langevin-MARINA, is inspired from an optimization algorithm called MARINA (Gorbunov et al., 2021) for which we give a new convergence proof in Appendix B. This new proof draws connections between optimization (MARINA) and sampling (Langevin-MARINA).

1.3 Paper organization

In Section 2, we review some background material on sampling and optimization. Next, we introduce our federated learning setup in Section 3. In Section 4 we give our main algorithm, Langevin-MARINA, and our main complexity results. We conclude in Section 5. The proofs, including the new proofs of existing results for MARINA, are deferred to the Appendix.

2 Preliminaries

2.1 Mathematical problem

Throughout this paper, we consider a nonconvex function $F : \mathbb{R}^d \to \mathbb{R}$ which is assumed to be L-smooth, i.e., differentiable with an L-Lipschitz continuous gradient: $\|\nabla F(x) - \nabla F(y)\| \leq L \|x - y\|$. As often in machine learning, F can be seen as a training loss and takes the form of a finite sum

$$F = \sum_{i=1}^{n} F_i, \tag{1}$$

where each F_i can be seen as the loss associated to the dataset stored in the device i. Assuming that $\int \exp(-F(x))dx \in (0, \infty)$, we denote by

$$\pi \propto \exp(-F), \tag{2}$$

the probability distribution whose density is proportional to $\exp(-F)$. We take a Bayesian approach: instead of minimizing F, our goal is to generate random samples from π, which can be seen as the posterior distribution of some Bayesian model.

2.2 Sampling and Optimal Transport

We denote by $\mathcal{P}_2(\mathbb{R}^d)$ the set of Borel measures σ on $\mathbb{R}^d$ with finite second moment, that is $\int \|x\|^2 \, d\sigma(x) < +\infty$. For every $\sigma, \nu \in \mathcal{P}_2(\mathbb{R}^d)$, $\Gamma(\sigma, \nu)$ is the set of all the coupling measures between σ and ν on $\mathbb{R}^d \times \mathbb{R}^d$, that is $\gamma \in \Gamma(\sigma, \nu)$ if and only if $\sigma(dx) = \gamma(dx, \mathbb{R}^d)$ and $\nu(dy) = \gamma(\mathbb{R}^d, dy)$. The Wasserstein distance between σ and ν is defined by

$$W_2(\sigma, \nu) = \sqrt{\inf_{\gamma \in \Gamma(\sigma,\nu)} \int \|x - y\|^2 \, \gamma(dx, dy)}. \tag{3}$$

The Wasserstein distance $W_2(\cdot, \cdot)$ is a metric on $\mathcal{P}_2(\mathbb{R}^d)$ and the metric space $\big(\mathcal{P}_2(\mathbb{R}^d), W_2\big)$ is called the Wasserstein space, see Ambrosio et al. (2005).

The Kullback-Leibler (KL) divergence w.r.t. π can be seen as the map from the Wasserstein space to $(0, \infty]$ defined for every $\sigma \in \mathcal{P}_2(\mathbb{R}^d)$ by

$$D_{\mathrm{KL}}(\sigma \mid \pi) := \begin{cases} \int \log(\frac{\sigma}{\pi})(x) d\sigma(x) & \text{if } \sigma \ll \pi \\ \infty & \text{else,} \end{cases} \tag{4}$$

where $\sigma \ll \pi$ means that σ is absolutely continuous w.r.t. π and $\frac{\sigma}{\pi}$ denotes the density of σ w.r.t. π. The KL divergence is always nonnegative and it is equal to zero if and only if $\sigma = \pi$. Therefore, assuming that $\pi \in \mathcal{P}_2(\mathbb{R}^d)$, π can be seen as the solution to the optimization problem

$$\min_{\sigma \in \mathcal{P}_2(\mathbb{R}^d)} D_{\mathrm{KL}}\left(\sigma \mid \pi\right). \tag{5}$$

Viewing π as a minimizer of the KL divergence is the cornerstone of our approach. Indeed, we shall view the proposed algorithm as an optimization algorithm to solve Problem equation 5. In particular, following Wibisono (2018); Durmus et al. (2019), we shall view the Langevin algorithm as a first order algorithm over the Wasserstein space. In particular, the relative Fisher information, defined as

$$J_\pi\left(\sigma\right) := \begin{cases} \int \left\|\nabla \log(\frac{\sigma}{\pi})(x)\right\|^2 d\sigma(x) & \text{if } \sigma \ll \pi \\ \infty, & \text{else,} \end{cases} \tag{6}$$

will play the role of the squared norm of the gradient of the objective function. Indeed, by defining a differential structure over the Wasserstein space (Ambrosio et al., 2005), the relative Fisher information is the squared norm of the gradient of the KL divergence.

The analysis of nonconvex optimization algorithm minimizing F often relies on a gradient domination condition relating the squared norm of the gradient to the function values, such as the following Lojasiewicz condition (Blanchet & Bolte, 2018; Karimi et al., 2016)

$$F(x) - \min F \leq \frac{1}{2\mu}\|\nabla F(x)\|^2. \tag{7}$$

The latter condition has a well-known analogue over the Wasserstein space for the KL divergence, called logarithmic Sobolev inequality.

Assumption 2.1 (Logarithmic Sobolev Inequality (LSI))**.** The distribution π satisfies the logarithmic Sobolev inequality with constant μ: for all $\sigma \in \mathcal{P}_2(\mathbb{R}^d)$,

$$D_{\mathrm{KL}}\left(\sigma \mid \pi\right) \leq \frac{1}{2\mu} J_\pi\left(\sigma\right). \tag{8}$$

The log Sobolev inequality has been studied in the optimal transport community (Villani, 2009, Chapter 21) and holds for several target distributions π. First, if F is μ-strongly convex (we say that π is strongly log-concave), then LSI holds with constant μ. Besides, LSI is preserved under bounded perturbations and Lipschitz mappings (Vempala & Wibisono, 2019, Lemma 16, 19). Therefore, small perturbations of strongly-log concave distributions satisfy LSI. For instance, if we add to a Gaussian distribution another Gaussian distribution with a small weight, then the resulting mixture of Gaussian distributions is not log-concave but it satisfies LSI. One can also find compactly supported examples, see Wibisono (2019, Introduction).

LSI is a condition that has been used in the analysis of the Langevin algorithm

$$x_{k+1} = x_k - h\nabla F(x_k) + \sqrt{2h} Z_{k+1}, \tag{9}$$

where $h > 0$ is a step size and (Z_k) a sequence of i.i.d standard Gaussian vectors over $\mathbb{R}^d$. Langevin algorithm can be seen as a gradient descent algorithm for F to which a Gaussian noise is added at each step. For instance, Wibisono (2019) showed that the distributions ρ_k of x_k converges rapidly towards the target distribution π in terms of the KL divergence under LSI. Other metrics have also been considered, such a the Total Variation distance

$$\|\sigma - \pi\|_{TV} = \sup_{A \in B(\mathbb{R}^d)} |\sigma(A) - \pi(A)|, \tag{10}$$

where $B(\mathbb{R}^d)$ denotes the Borel sigma field of $\mathbb{R}^d$.

Algorithm 1 MARINA (Gorbunov et al., 2021)

Input: Starting point x_0, step-size h, number of iterations K
for $i = 1, 2, \cdots, n$ in parallel **do**
  Device i computes MARINA estimator g_0^i
  Device i uploads g_0^i to the central server
end for
Server aggregates $g_0 = \frac{1}{n}\sum_{i=1}^n g_0^i$
for $k = 0, 1, 2, \cdots, K-1$ **do**
  Server broadcasts g_k to all devices i
  for $i = 1, 2, \cdots, n$ in parallel **do**
    Device i performs $x_{k+1} = x_k - hg_k$
    Device i computes MARINA estiamtor g_{k+1}^i
    Device i uploads g_{k+1}^i to the central server
  end for
  Server aggregates $g_{k+1} = \frac{1}{n}\sum_{i=1}^n g_{k+1}^i$
end for
Return: $\hat{x}^K$ chosen uniformly at random from $(x_k)_{k=0}^K$ or last point x_K

3 Federated learning

3.1 Example of a federated learning algorithm: MARINA

We now describe our federated learning setup. A central server and a number of devices are required to run a training algorithm in a distributed manner. Each device $i \in \{1, \ldots, n\}$ owns a dataset. A meta federated learning algorithm is as follows: at each iteration, each device performs some local computation (for instance, a gradient computation) using its own dataset, then each device sends the result of that computation to the central server. The central server aggregates the results and sends that aggregation back to the devices. Algorithm 1 provides an example of a federated learning algorithm. This algorithm, called MARINA, was introduced and analyzed in Gorbunov et al. (2021) under the assumption that the sequence (g_k) is a MARINA estimator of the gradient as defined below.

Definition 3.1. Given a sequence produced by an algorithm $(x_k)_k$, a MARINA estimator of the gradient is a random sequence $(g_k)_k$ satisfying $\mathrm{E}[g_k] = \mathrm{E}[\nabla F(x_k)]$ and

$$\mathbf{G}_{k+1} \leq (1-p)\mathbf{G}_k + (1-p)L^2\alpha\mathrm{E}\left[\|x_{k+1} - x_k\|^2\right] + \theta, \tag{11}$$

where $\mathbf{G}_k = \mathrm{E}\left[\|g_k - \nabla F(x_k)\|^2\right]$ and $0 < p \leq 1$, $\alpha, \theta \geq 0$.

MARINA iterations can be rewritten as

$$x_{k+1} = x_k - hg_k. \tag{12}$$

Since $\mathrm{E}[g_k] = \mathrm{E}[\nabla F(x_k)]$, MARINA can be seen as a SGD distributed between the central server and the devices, to minimize the nonconvex function F.

Recall that in federated learning, the main bottleneck is the communication cost between server and devices. The key aspect of MARINA is that it achieves the state of the art in terms of communication complexity among nonconvex optimization algorithms. The reason is the following: there are practical examples of MARINA estimators (g_k) that rely on a compression step before the communication and are therefore cheap to communicate[1].

[1]Note that we did not specify how g_k^i is computed in Algorithm 1. For the moment we only require g_k to be a MARINA estimator

3.2 Compression operator

MARINA estimators are typically cheaper to communicate than full gradient because they involve a compression step, making the communication light. A compression operator is a random map from $\mathbb{R}^d$ to $\mathbb{R}^d$ with the following properties.

Definition 3.2 (Compression)**.** A stochastic mapping $\mathcal{Q} : \mathbb{R}^d \to \mathbb{R}^d$ is a compression operator if there exists $\omega > 0$ such that for any $x \in \mathbb{R}^d$,

$$\mathrm{E}\left[\mathcal{Q}(x)\right] = x, \quad \mathrm{E}\left[\|\mathcal{Q}(x) - x\|^2\right] \leq \omega \|x\|^2. \tag{13}$$

The first equation states that $\mathcal{Q}$ is unbiased and the second equation states that the variance of the compression operator has a quadratic growth. There are lots of compression operators satisfying equation 13, explicit examples based on quantization and/or sparsification are given in Alistarh et al. (2017); Horváth et al. (2019a). In general, compressed quantities are cheap to communicate.

We now fix a compression operator and give examples of gradient estimators involving compression, which are provably MARINA estimators.

3.3 MARINA gradient estimators

In every examples presented below, $g_k := \frac{1}{n}\sum_{i=1}^n g_k^i$, where g_k^i is meant to approximate $\nabla F_i(x_k)$, i.e., the gradient at the device i, at step k. We consider any random sequence (x_k) generated by a training algorithm.

3.3.1 Vanilla MARINA gradient estimator

The vanilla MARINA gradient estimator is defined as the average $g_k := \frac{1}{n}\sum_{i=1}^n g_k^i$, where for every $i \in \{1, \ldots, n\}$, $g_0^i = \nabla F_i(x_0)$ and

$$g_{k+1}^i := \begin{cases} \nabla F_i(x_{k+1}) & \text{prob. } p > 0 \\ g_k + \mathcal{Q}\left(\nabla F_i(x_{k+1}) - \nabla F_i(x_k)\right) & \text{prob. } \ 1-p \end{cases}. \tag{14}$$

Node i randomly (and independently of $\mathcal{Q}, g_k, x_{k+1}$) computes and uploads to the server the gradient of F_i at point x_{k+1} with probability p. Else, node i computes and uploads the compressed difference of the gradients of F_i between points x_{k+1} and x_k: $\mathcal{Q}\left(\nabla F_i(x_{k+1}) - \nabla F_i(x_k)\right)$. Note that the server already knows g_k from the previous iteration, so no need to send g_k to the server. In average, device i computes a full gradient every $1/p$ steps (p is very small, see Remark 3.6 below).

Proposition 3.3. *Assume that F_i is L_i-smooth for all $i \in \{1, \ldots, n\}$, i.e.,*

$$\|\nabla F_i(x) - \nabla F_i(y)\| \leq L_i \|x - y\|, \quad \forall x, y \in \mathbb{R}^d. \tag{15}$$

Then, $g_k = \frac{1}{n}\sum_{i=1}^n g_k^i$, where g_k^i is defined by equation 14 is a MARINA *estimator in the sense of Definition 3.1, with $\alpha = \frac{\omega \sum_{i=1}^n L_i^2}{n^2 L^2}, \theta = 0$.*

This proposition means that the MARINA estimator reduces the variance induced by the compression operator. Note that, if F_i is L_i-smooth, then $L \leq \frac{1}{n}\sum_{i=1}^n L_i$.

3.3.2 Finite sum case

Consider the finite sum case where each F_i is a sum over the data points stored in the device i: $F_i = \sum_{j=1}^N F_{ij}$, where F_{ij} is the loss associated to the data point j in the device i. The finite sum MARINA gradient estimator is analogue to the vanilla estimator, but with subsampling. It is defined as the average $g_k := \frac{1}{n}\sum_{i=1}^n g_k^i$, where for every $i \in \{1, \ldots, n\}$, $g_0^i = \nabla F_i(x_0)$ and

$$g_{k+1}^i := \begin{cases} \nabla F_i(x_{k+1}) & \text{prob. } p > 0 \\ g_k + \mathcal{Q}\left(\frac{1}{b'}\sum_{j \in I'_{i,k}} \left(\nabla F_{ij}(x_{k+1}) - \nabla F_{ij}(x_k)\right)\right) & \text{prob. } 1-p \end{cases}, \tag{16}$$

where b' is the minibatch size and $I'_{i,k}$ is the set of the indices in the minibatch, $\left|I'_{i,k}\right| = b'$. Here, $(I'_{i,k})_{i,k}$ are i.i.d random sets consisting in b' i.i.d samples from the uniform distribution over $\{1, \ldots, N\}$.

Proposition 3.4. *Assume that F_{ij} is L_{ij}-smooth for all $i \in \{1, \ldots, n\}$ and all $j \in \{1, \ldots, N\}$. Then, $g_k = \frac{1}{n}\sum_{i=1}^n g_k^i$, where g_k^i is defined by equation 16 is a* MARINA *estimator in the sense of Definition 3.1, with $\alpha = \frac{\omega \sum_{i=1}^n L_i^2 + (1+\omega)\frac{\sum_{i=1}^n \mathcal{L}_i^2}{b'}}{n^2 L^2}$, where $\mathcal{L}_i = \max_{j\in[N]} L_{ij}$ and $\theta = 0$.*

This proposition means that the MARINA estimator reduces the variance induced by the compression operator and the variance induced by the subsampling.

3.3.3 Online case

Consider the online case where each F_i is an expectation over the randomness of a stream of data arriving online: $F_i(x) = \mathbb{E}_{\xi_i \sim \mathcal{D}_i}\left[F_{\xi_i}(x)\right]$, where F_{ξ_i} is the loss associated to a data point ξ_i in device i. The online MARINA gradient estimator is analogue to the finite sum estimator, but with subsampling from $\mathcal{D}_i$ instead of the uniform over $\{1, \ldots, N\}$. Moreover, the full gradient ∇F_i is never computed (because it is intractable in online learning). The online MARINA gradient estimator is defined as the average $g_k := \frac{1}{n}\sum_{i=1}^n g_k^i$, where for every $i \in \{1, \ldots, n\}$, $g_0^i = \frac{1}{b}\sum_{\xi \in I_{i,0}} \nabla F_\xi\left(x_0\right)$ and

$$g_{k+1}^i := \begin{cases} \frac{1}{b}\sum_{\xi \in I_{i,k}} \nabla F_\xi\left(x_{k+1}\right) & \text{prob. } p > 0 \\ g_k + \mathcal{Q}\left(\frac{1}{b'}\sum_{\xi \in I'_{i,k}} \left(\nabla F_\xi\left(x_{k+1}\right) - \nabla F_\xi\left(x_k\right)\right)\right) & \text{prob. } 1-p \end{cases}, \tag{17}$$

where $I'_{i,k}$ and $I_{i,k}$ are the set of the indices in a minibatch, $\left|I'_{i,k}\right| = b'$ and $|I_{i,k}| = b$. Here, $(I_{i,k})_{i,k}$ (resp. $(I'_{i,k})_{i,k}$) is an i.i.d random set consisting in b (resp. b') i.i.d samples from $\mathcal{D}_i$. Besides, in the online case only, we make a bounded gradient assumption.

Proposition 3.5. *Assume that for all $i \in \{1, \ldots, n\}$ there exists $\sigma_i \geq 0$ such that for all $x \in \mathbb{R}^d$,*

$$\begin{aligned} &\mathbb{E}_{\xi_i \sim \mathcal{D}_i}\left[\nabla F_{\xi_i}(x)\right] = \nabla F_i(x), \\ &\mathbb{E}_{\xi_i \sim \mathcal{D}_i}\left[\left\|\nabla F_{\xi_i}(x) - \nabla F_i(x)\right\|^2\right] \leq \sigma_i^2. \end{aligned}$$

Moreover, assume that F_{ξ_i} is L_{ξ_i}-smooth $\mathcal{D}_i$ almost surely. Then, $g_k = \frac{1}{n}\sum_{i=1}^n g_k^i$, where g_k^i is defined by equation 17 is a MARINA *estimator in the sense of Definition 3.1, with $\alpha = \frac{\omega \sum_{i=1}^n L_i^2 + (1+\omega)\frac{\sum_{i=1}^n \mathcal{L}_i^2}{b'}}{n^2 L^2}$, where $\mathcal{L}_i$ is the L^∞ norm of L_{ξ_i} (where $\xi_i \sim \mathcal{D}_i$) and $\theta = \frac{p\sum_{i=1}^n \sigma_i^2}{n^2 b}$.*

Note that in the online case $\theta \neq 0$ in general, because the online MARINA estimator does not reduce to zero the variance induced by the subsampling. That is because the loss F_i is an expectation (and not a finite sum *a priori*), for which variance reduction to zero is impossible in general. However, θ can be made small by taking a large minibatch size b or a small probability p. Moreover, in the online case, $\mathbf{G}_0 \neq 0$ unlike in the two other cases.

Remark 3.6. Typically, p is very small, for instance, Gorbunov et al. (2021) chooses $p = \zeta_{\mathcal{Q}}/d$ in equation 14, $p = \min\left\{\zeta_{\mathcal{Q}}/d, b'/(N+b')\right\}$ in equation 16 and $p = \min\left\{\zeta_{\mathcal{Q}}/d, b'/(b+b')\right\}$ in equation 17, where $\zeta_{\mathcal{Q}} = \sup_{x\in\mathbb{R}^d} \mathbb{E}\left[\|\mathcal{Q}(x)\|_0\right]$ and $\|y\|_0$ is the number of non-zero components of $y \in \mathbb{R}^d$. Therefore, with high probability $1-p$, the compressed difference $\mathcal{Q}\left(\nabla F_i(x_{k+1}) - \nabla F_i(x_k)\right)$ is sent to the server. Sending compressed quantities has a low communication cost (Alistarh et al., 2017; Horváth et al., 2019a).

Remark 3.7. Propositions 3.3, 3.4 and 3.5 can be found in Gorbunov et al. (2021, Equation 21, 33, 46) in the case where (x_k) is a stochastic gradient descent algorithm (i.e., $x_{k+1} = x_k - hg^k$). In the general case where (x_k) is produced by any algorithm, we found out that the proofs of these Propositions are the same, but we reproduce these proofs in the Appendix for the sake of completeness.

4 Langevin-MARINA

In this section, we give our main algorithm, Langevin-MARINA, and study its convergence for sampling from $\pi \propto \exp(-F)$. We prove convergence bounds in KL divergence, 2-Wasserstein distance and Total Variation distance under LSI.

4.1 Proposed algorithm

Our main algorithm, Langevin-MARINA can be seen as a Langevin variant of MARINA which adds a Gaussian noise at each step of MARINA (Algorithm 1). Our motivation is to obtain a Langevin algorithm whose communication complexity is similar to that of MARINA. Alternatively, one can see Langevin-MARINA as a MARINA variant of Langevin algorithm which uses a MARINA estimator of the gradient $g_k = \frac{1}{n}\sum_{i=1}^n g_k^i$ instead of the full gradient $\nabla F(x_k)$ in the Langevin algorithm:

$$x_{k+1} = x_k - hg_k + \sqrt{2h}Z_{k+1}. \tag{18}$$

Langevin-MARINA is presented in Algorithm 2.

Algorithm 2 Langevin-MARINA (proposed algorithm)

Input: Starting point $x_0 \sim \rho_0$, step-size h, number of iterations K
for $i = 1, 2, \cdots, n$ in parallel **do**
    Device i computes MARINA estimator g_0^i
    Device i uploads g_0^i to the central server
end for
Server aggregates $g_0 = \frac{1}{n}\sum_{i=1}^n g_0^i$
for $k = 0, 1, 2, \cdots, K-1$ **do**
    Server broadcasts g_k to all devices i
    Server draws a Gaussian vector $Z_{k+1} \sim \mathcal{N}(0, I_d)$
    Server performs $x_{k+1} = x_k - hg_k + \sqrt{2h}Z_{k+1}$
    Server broadcasts x_{k+1} to all devices i
    for $i = 1, 2, \cdots, n$ in parallel **do**
        Device i computes MARINA estimator g_{k+1}^i
        Device i uploads g_{k+1}^i to the central server
    end for
    Server aggregates $g_{k+1} = \frac{1}{n}\sum_{i=1}^n g_{k+1}^i$
end for
Return: x_K

4.2 Communication complexity of Langevin-MARINA

Compared to MARINA where n equivalent stochastic gradient descent steps are performed by the devices, here the Langevin step is performed only once, by the server. This allows for computation savings (n times less computations) at the cost of more communication: the iterates x_k need to be broadcast by the server to the devices. Therefore, the communication complexity of the sampling algorithm Langevin-MARINA is higher than the communication complexity of the optimization algorithm MARINA.

However, compared to the concurrent sampling algorithm of Vono et al. (2022), the communication complexity per iteration of Langevin-MARINA is equivalent. Comparing the communication complexity of Langevin-MARINA to that of FA-LD (Deng et al., 2021) is more difficult because FA-LD makes communication savings by performing communication rounds only after a number T of local updates, instead of using a compression operator after each local update like Langevin-MARINA.

4.3 Convergence theory for Langevin-MARINA

We first prove the convergence of Langevin-MARINA in KL-divergence.

Theorem 4.1. *Assume that LSI (Assumption 2.1) holds and that $g_k = \frac{1}{n}\sum_{i=1}^n g_k^i$ is a* MARINA *estimator in the sense of Definition 3.1. If*

$$0 < h \leq \min\left\{\frac{1}{14L}\sqrt{\frac{p}{1+\alpha}}, \frac{p}{6\mu}\right\}, \tag{19}$$

then

$$D_{\mathrm{KL}}\left(\rho_K \mid \pi\right) \leq e^{-\mu K h}\Psi_3 + \frac{1-e^{-K\mu h}}{\mu}\tau, \tag{20}$$

where $\Psi_3 = D_{\mathrm{KL}}\left(\rho_0 \mid \pi\right) + \frac{1-e^{-\mu h}}{\mu}C\boldsymbol{G}_0, \tau = \left(2L^2 + C(1-p)L^2\alpha\right)\left(8Lh^2d + 4dh\right) + C\theta, C = \frac{8L^2h^2\beta+2\beta}{1-(1-p)(4L^2h^2\alpha+1)\beta}, \beta = e^{\mu h}$ *and ρ_k is the distribution of x_k for every k.*

In particular, if $\theta = 0$, set $h = \mathcal{O}\left(\frac{\mu p\varepsilon}{L^2(1+\alpha)d}\right)$, $K = \Omega\left(\frac{L^2(1+\alpha)d}{\mu^2 p\varepsilon}\log\left(\frac{\Psi_3}{\varepsilon}\right)\right)$ then $D_{\mathrm{KL}}\left(\rho_K \mid \pi\right) \leq \varepsilon$. If $\theta \neq 0$, there will always be an extra residual term $\frac{1-e^{-K\mu h}}{\mu}C\theta$ in the right hand side of equation 20 which cannot be diminished to 0 by setting K and h. However in the online case (Section 3.3.3), $\theta = \frac{p\sum_{i=1}^n \sigma_i^2}{n^2 b}$, we can set $b = \Omega\left(\frac{\sum_{i=1}^n \sigma_i^2}{\mu n^2 \varepsilon}\right)$ to make the residual term small $\frac{1-e^{-K\mu h}}{\mu}C\theta = \mathcal{O}(\varepsilon)$.

From this theorem, we can obtain complexity results in Total Variation distance and 2-Wasserstein distance. Indeed, Pinsker inequality states that

$$\|\sigma - \pi\|_{TV} \leq \sqrt{\frac{1}{2}H_\pi(\sigma)}, \forall \sigma, \nu \in \mathcal{P}_2(\mathbb{R}^d), \tag{21}$$

and LSI implies Talagrand's T_2 inequality

$$W_2^2(\sigma, \pi) \leq \frac{2}{\mu}D_{\mathrm{KL}}\left(\sigma \mid \pi\right), \forall \sigma \in \mathcal{P}_2(\mathbb{R}^d), \tag{22}$$

see Villani (2009, Chapter 21).

Corollary 4.2. *Let assumptions and parameters be as in Theorem 4.1. Then,*

$$\|\rho_K - \pi\|_{TV}^2 \leq \frac{1}{2}\left(e^{-\mu K h}\Psi_3 + \frac{1-e^{-K\mu h}}{\mu}\tau\right), \tag{23}$$

and

$$W_2^2(\rho_K, \pi) \leq \frac{2}{\mu}\left(e^{-\mu K h}\Psi_3 + \frac{1-e^{-K\mu h}}{\mu}\tau\right). \tag{24}$$

In particular, if $\theta = 0$, set $h = \mathcal{O}\left(\frac{\mu p\varepsilon^2}{L^2(1+\alpha)d}\right)$, $K = \Omega\left(\frac{L^2(1+\alpha)d}{\mu^2 p\varepsilon^2}\log\left(\frac{\Psi_3}{\varepsilon^2}\right)\right)$ then $\|\rho_K - \pi\|_{TV} \leq \varepsilon$. If $\theta \neq 0$, there will always be an extra residual term $\frac{1-e^{-K\mu h}}{\mu}C\theta$ in the right hand side of equation 23 which cannot be diminished to 0 by setting K and h. However in the online case (Section 3.3.3), $\theta = \frac{p\sum_{i=1}^n \sigma_i^2}{n^2 b}$, we can set $b = \Omega\left(\frac{\sum_{i=1}^n \sigma_i^2}{\mu n^2 \varepsilon^2}\right)$ to make the residual term small $\frac{1-e^{-K\mu h}}{\mu}C\theta = \mathcal{O}(\varepsilon^2)$.

Finally, if $\theta = 0$, set $h = \mathcal{O}\left(\frac{\mu^2 p\varepsilon^2}{L^2(1+\alpha)d}\right)$, $K = \Omega\left(\frac{L^2(1+\alpha)d}{\mu^3 p\varepsilon^2}\log\left(\frac{\Psi_3}{\mu\varepsilon^2}\right)\right)$ then $W_2(\rho_K, \pi) \leq \varepsilon$. If $\theta \neq 0$, there will always be an extra residual term $\frac{1-e^{-K\mu h}}{\mu^2}C\theta$ in the right hand side of equation 24 which cannot be diminished to 0 by setting K and h. However in the online case, $\theta = \frac{p\sum_{i=1}^n \sigma_i^2}{n^2 b}$, we can set $b = \Omega\left(\frac{\sum_{i=1}^n \sigma_i^2}{\mu^2 n^2 \varepsilon^2}\right)$ to make the residual term small $\frac{1-e^{-K\mu h}}{\mu^2}C\theta = \mathcal{O}(\varepsilon^2)$.

We can compare our results to Vono et al. (2022); Deng et al. (2021) which also obtain bounds in 2-Wasserstein distance under the stronger assumption that F is μ-strongly convex, see Table 1. In Vono et al.

(2022), they need $h = \mathcal{O}\left(\frac{\varepsilon^2}{dl}\right)$ and $K = \Omega\left(\frac{dl}{\varepsilon^2}\log\left(\frac{W_2^2(\rho_0,\pi)}{\varepsilon^2}\right)\right)$, where they compute the full gradient every l step (we do not include the dependence in μ, L in their result for the sake of simplicity). One can think of l as $1/p$. In Deng et al. (2021), they require $h = \mathcal{O}\left(\frac{\mu^2\varepsilon^2}{dL^2(T^2\mu+L)}\right)$ and $K = \Omega\left(\frac{dL^2(T^2\mu+L)}{\mu^3\varepsilon^2}\log\left(\frac{d}{\varepsilon^2}\right)\right)$ to get $W_2(\rho_K,\pi) \leq \varepsilon$, where T denotes the number of local updates.

5 Conclusion

We introduced a communication efficient variant of Langevin algorithm for federated learning called Langevin-MARINA. We studied the complexity of this algorithm in terms of KL divergence, Total Variation distance and 2-Wasserstein distance. Unlike existing works on sampling for federated learning, we only require the target distribution to satisfy LSI, which allows the target distribution to not be log-concave.

Langevin-MARINA is inspired by a optimization algorithm for federated learning called MARINA. MARINA achieves the state of the art in communication complexity among nonconvex optimization algorithms. However, Langevin-MARINA requires more communication rounds than MARINA. The fundamental reason for this is that Langevin-MARINA needs to communicate the Gaussian noise. To solve this issue, one approach is to design a Langevin algorithm allowing for the compression of the Gaussian noise. Another approach is to put the same random seed in each device in order to ensure that they all generate the same Z_{k+1} at step k. This approach does not require the communication of the Gaussian noise.

Finally, simple experiments have shown similar performance for Langevin-Marina and QLSD^{++} Vono et al. (2022), although the convergence of QLSD^{++} under LSI has not been established. Therefore, we wonder if a convergence theory of QLSD^{++} under LSI would be possible. We leave this question for future work.

References

Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. Qsgd: communication-efficient SGD via gradient quantization and encoding. *Advances in Neural Information Processing Systems*, 30, 2017.

Herbert Amann. *Ordinary differential equations: an introduction to nonlinear analysis*, volume 13. Walter de Gruyter, 2011.

Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2005.

Krishnakumar Balasubramanian, Sinho Chewi, Murat A Erdogdu, Adil Salim, and Matthew Zhang. Towards a theory of non-log-concave sampling: first-order stationarity guarantees for Langevin Monte Carlo. *arXiv preprint arXiv:2202.05214*, 2022.

Adrien Blanchet and Jérôme Bolte. A family of functional inequalities: Łojasiewicz inequalities and displacement convex functions. *Journal of Functional Analysis*, 275(7):1650–1673, 2018.

Xiang Cheng, Niladri S Chatterji, Yasin Abbasi-Yadkori, Peter L Bartlett, and Michael I Jordan. Sharp convergence rates for Langevin dynamics in the nonconvex setting. *arXiv preprint arXiv:1805.01648*, 2018.

Sinho Chewi, Murat A Erdogdu, Mufan Bill Li, Ruoqi Shen, and Matthew Zhang. Analysis of Langevin Monte Carlo from Poincaré to log-Sobolev. *arXiv preprint arXiv:2112.12662*, 2021.

Arnak 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.

Rudrajit Das, Anish Acharya, Abolfazl Hashemi, Sujay Sanghavi, Inderjit S Dhillon, and Ufuk Topcu. Faster non-convex federated learning via global and local momentum. *arXiv preprint arXiv:2012.04061*, 2020.

Wei Deng, Yi-An Ma, Zhao Song, Qian Zhang, and Guang Lin. On convergence of federated averaging Langevin dynamics. *arXiv preprint arXiv:2112.05120*, 2021.

Dimitrios Dimitriadis, Mirian Hipolito Garcia, Daniel Madrigal Diaz, Andre Manoel, and Robert Sim. Flute: A scalable, extensible framework for high-performance federated learning simulations. *arXiv preprint arXiv:2203.13789*, 2022.

Alain Durmus and Eric Moulines. Nonasymptotic convergence analysis for the unadjusted Langevin algorithm. *The Annals of Applied Probability*, 27(3):1551–1587, 2017.

Alain Durmus, Szymon Majewski, and Błażej Miasojedow. Analysis of Langevin Monte Carlo via convex optimization. *The Journal of Machine Learning Research*, 20(1):2666–2711, 2019.

Khaoula El Mekkaoui, Diego Mesquita, Paul Blomstedt, and Samuel Kaski. Federated stochastic gradient Langevin dynamics. In *Uncertainty in Artificial Intelligence*, pp. 1703–1712. PMLR, 2021.

Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.

Eduard Gorbunov, Konstantin P Burlachenko, Zhize Li, and Peter Richtárik. Marina: faster non-convex distributed learning with compression. In *International Conference on Machine Learning*, pp. 3788–3798. PMLR, 2021.

Farzin Haddadpour, Mohammad Mahdi Kamani, Aryan Mokhtari, and Mehrdad Mahdavi. Federated learning with compression: unified analysis and sharp guarantees. In *International Conference on Artificial Intelligence and Statistics*, pp. 2350–2358. PMLR, 2021.

Samuel Horváth, Chen-Yu Ho, Ludovit Horvath, Atal Narayan Sahu, Marco Canini, and Peter Richtárik. Natural compression for distributed deep learning. *arXiv preprint arXiv:1905.10988*, 2019a.

Samuel Horváth, Dmitry Kovalev, Konstantin Mishchenko, Sebastian Stich, and Peter Richtárik. Stochastic distributed learning with gradient quantization and variance reduction. *arXiv preprint arXiv:1904.05115*, 2019b.

Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021.

Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the Polyak-Lojasiewicz condition. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 795–811. Springer, 2016.

Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, pp. 5132–5143. PMLR, 2020.

Jakub Konečnỳ, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016.

Yi-An Ma, Niladri Chatterji, Xiang Cheng, Nicolas Flammarion, Peter Bartlett, and Michael I Jordan. Is there an analog of Nesterov acceleration for MCMC? *arXiv preprint arXiv:1902.00996*, 2019.

Mateusz B Majka, Aleksandar Mijatović, and Łukasz Szpruch. Nonasymptotic bounds for sampling algorithms without log-concavity. *The Annals of Applied Probability*, 30(4):1534–1581, 2020.

Jonathan C Mattingly, Andrew M Stuart, and Desmond J Higham. Ergodicity for SDEs and approximations: locally Lipschitz vector fields and degenerate noise. *Stochastic processes and their applications*, 101(2): 185–232, 2002.

Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pp. 1273–1282. PMLR, 2017.

Lam M Nguyen, Jie Liu, Katya Scheinberg, and Martin Takáč. Sarah: a novel method for machine learning problems using stochastic recursive gradient. In *International Conference on Machine Learning*, pp. 2613–2621. PMLR, 2017.

Santosh Vempala and Andre Wibisono. Rapid convergence of the unadjusted Langevin algorithm: isoperimetry suffices. *Advances in Neural Information Processing Systems*, 32, 2019.

Cédric Villani. *Optimal transport: old and new*, volume 338. Springer, 2009.

Maxime Vono, Vincent Plassier, Alain Durmus, Aymeric Dieuleveut, and Eric Moulines. Qlsd: quantised Langevin stochastic dynamics for Bayesian federated learning. In *International Conference on Artificial Intelligence and Statistics*, pp. 6459–6500. PMLR, 2022.

Andre Wibisono. Sampling as optimization in the space of measures: The Langevin dynamics as a composite optimization problem. In *Conference on Learning Theory (COLT)*, pp. 2093–3027, 2018.

Andre Wibisono. Proximal Langevin algorithm: rapid convergence under isoperimetry. *arXiv preprint arXiv:1911.01469*, 2019.