
Faster First-Order Methods for Stochastic Non-Convex Optimization on Riemannian Manifolds

Pan Zhou*

*National University of Singapore
pzhou@u.nus.edu

Xiao-Tong Yuan†

†Nanjing University of Information Science & Technology
xyuan@nuist.edu.cn

Jiashi Feng*

elefjia@nus.edu.sg

Abstract

SPIDER (Stochastic Path Integrated Differential Estimator) is an efficient gradient estimation technique developed for non-convex stochastic optimization. Although having been shown to attain nearly optimal computational complexity bounds, the SPIDER-type methods are limited to linear metric spaces. In this paper, we introduce the Riemannian SPIDER (R-SPIDER) method as a novel nonlinear-metric extension of SPIDER for efficient non-convex optimization on Riemannian manifolds. We prove that for finite-sum problems with n components, R-SPIDER converges to an ϵ -accuracy stationary point within $\mathcal{O}(\min(n + \frac{\sqrt{n}}{\epsilon^2}, \frac{1}{\epsilon^3}))$ stochastic gradient evaluations, which is sharper in magnitude than the prior Riemannian first-order methods. For online optimization, R-SPIDER is shown to converge with $\mathcal{O}(\frac{1}{\epsilon^3})$ complexity which is, to the best of our knowledge, the first non-asymptotic result for online Riemannian optimization. Especially, for gradient dominated functions, we further develop a variant of R-SPIDER and prove its linear convergence rate. Numerical results testify the computational efficiency of the proposed methods.

1 Introduction

We consider the following *finite-sum* and *online* non-convex problems on a Riemannian manifold \mathcal{M} :

$$\min_{\mathbf{x} \in \mathcal{M}} f(\mathbf{x}) := \begin{cases} \frac{1}{n} \sum_{i=1}^n f_i(\mathbf{x}) & \text{(finite-sum)} \\ \mathbb{E}[f(\mathbf{x}; \pi)] & \text{(online)} \end{cases}, \quad (1)$$

Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS) 2019, Naha, Okinawa, Japan. PMLR: Volume 89. Copyright 2019 by the author(s).

where $f : \mathcal{M} \mapsto \mathbb{R}$ is a smooth non-convex loss function. For the finite-sum problem, each individual loss $f_i(\mathbf{x})$ is associated with the i -th sample, while in online setting, the stochastic component $f(\mathbf{x}; \pi)$ is indexed by a random variable π . Such a formulation encapsulates several important finite-sum problems and their corresponding online counterparts, including principle component analysis (PCA) [1], low-rank matrix/tensor completion/recovery [2, 3, 4, 5, 6, 7, 8], dictionary learning [9, 10, 11, 12], Gaussian mixture models [13] and low-rank multivariate regression [14], to name a few.

One classic approach for solving problem (1) (or its convex counterpart) is to take it as a constrained optimization problem in ambient Euclidean space and find the minimizers via projected (stochastic) gradient descent [17, 18, 19]. This kind of methods, however, tend to suffer from high computational cost as projection onto certain manifolds (e.g., positive-definite matrices) could be expensive in large-scale learning problems [16].

As an appealing alternative, the Riemannian optimization methods have recently gained wide attention in machine learning [15, 16, 20, 21, 22, 23, 24]. In contrast to the Euclidean-projection based methods, the Riemannian methods directly move the iteration along the geodesic path towards the optimal solution, and thus can better respect the geometric structure of the problem in hand. Specifically, the Riemannian gradient methods have the following recursive form:

$$\mathbf{x}_{k+1} = \text{Exp}_{\mathbf{x}_k}(-\eta_k \mathbf{g}_k), \quad (2)$$

where \mathbf{g}_k is the gradient estimate of the full Riemannian gradient $\nabla f(\mathbf{x}_k)$, η_k denotes the learning rate, and the exponential mapping $\text{Exp}_{\mathbf{x}}(\mathbf{y})$, as defined in Section 2, maps \mathbf{y} in the tangent space at \mathbf{x} to $\text{Exp}_{\mathbf{x}}(\mathbf{y})$ on the manifold \mathcal{M} along a proper geodesic curve. For instance, Riemannian gradient descent (R-GD) uses the full Riemannian gradient $\mathbf{g}_k = \nabla f(\mathbf{x}_k)$ in Eqn. (2) and has been shown to have sublinear rate of convergence in geodesically convex problems [21]. To boost efficiency, Liu *et al.* [25] and Zhang *et al.* [20] further

Table 1: Comparison of IFO complexity for different Riemannian first-order stochastic optimization algorithms on the nonconvex problem (1) under finite-sum and online settings. The ϵ -accuracy solution is measured by the expected gradient norm $\mathbb{E}[\|\nabla f(\mathbf{x})\|] \leq \epsilon$. Here L , σ and ζ respectively denote the gradient Lipschitz constant, the gradient variance and the curvature parameter of the Riemannian manifold (see Section 2).

		Non-convex Problem	
		general non-convex	τ -gradient dominated
Finite-sum	R-SRG [15]	$\mathcal{O}\left(n + \frac{L^2}{\epsilon^4}\right)$	$\mathcal{O}\left((n + \tau^2 L^2) \log\left(\frac{1}{\epsilon}\right)\right)$
	R-SVRG [16]	$\mathcal{O}\left(n + \frac{\zeta n^{\frac{2}{3}}}{\epsilon^{\frac{2}{3}}}\right)$	$\mathcal{O}\left((n + \tau L \zeta^{\frac{1}{2}} n^{\frac{2}{3}}) \log\left(\frac{1}{\epsilon}\right)\right)$
	this work	$\mathcal{O}\left(\min\left(n + \frac{L\sqrt{n}}{\epsilon^2}, \frac{L\sigma}{\epsilon^3}\right)\right)$	$\mathcal{O}\left(\min\left((n + \tau L\sqrt{n}) \log\left(\frac{1}{\epsilon}\right), \frac{\tau L\sigma}{\epsilon}\right)\right)$
Online	this work	$\mathcal{O}\left(\frac{L\sigma}{\epsilon^3}\right)$	$\mathcal{O}\left(\frac{\tau L\sigma}{\epsilon}\right)$

introduced the Nesterov acceleration techniques [26] into R-GD with convergence rate significantly improved for geodesically convex functions.

To avoid the time-consuming full gradient computation required in R-GD, Riemannian stochastic optimization algorithms [15, 16, 22, 23, 24] leverage the decomposable (finite-sum) structure of problem (1). For instance, Bonnabel *et al.* [22] proposed R-SGD that only evaluates gradient of one (or a mini-batch) randomly selected sample for variable update per iteration. Though with good iteration efficiency, R-SGD converges slowly as it uses decaying learning rate for convergence guarantee due to its gradient variance. To tackle this issue, Riemannian stochastic variance-reduced gradient (R-SVRG) algorithms [16, 24] adapt SVRG [27] to problem (1). Benefiting from the variance-reduced technique, R-SVRG converges more stably and efficiently than R-SGD. More recently, inspired by the variance-reduced stochastic recursive gradient approach [28, 29], the Riemannian stochastic recursive gradient (R-SRG) algorithm [15] establishes a recursive equation to estimate the full Riemannian gradient so that the computational efficiency can be further improved (see Table 1).

SPIDER (Stochastic Path Integrated Differential Estimator) [30] is a recursive estimation method developed for tracking the history full gradients with significantly reduced computational cost. By combining SPIDER with normalized gradient methods, nearly optimal iteration complexity bounds can be attained for non-convex optimization in Euclidean space [30, 31]. Though appealing in vector space problems, it has not been explored for non-convex optimization in nonlinear metric spaces such as Riemannian manifold.

In this paper, we introduce the *Riemannian Stochastic Path Integrated Differential Estimator* (R-SPIDER) as a simple yet efficient extension of the SPIDER from Euclidean space to Riemannian manifolds. Specifically, for a proper positive integer p , at each time instance k with $\text{mod}(k, p) \equiv 0$, R-SPIDER first samples a large data batch \mathcal{S}_1 and estimates the initial full Riemannian gra-

dient $\nabla f(\mathbf{x}_k)$ as $\mathbf{v}_k = \nabla f_{\mathcal{S}_1}(\mathbf{x}_k) = \frac{1}{|\mathcal{S}_1|} \sum_{i \in \mathcal{S}_1} f_i(\mathbf{x}_k)$. Then at each of the next $p - 1$ iterations, it samples a smaller mini-batch \mathcal{S}_2 and estimates/tracks $\nabla f(\mathbf{x}_k)$:

$$\mathbf{v}_k = \nabla f_{\mathcal{S}_2}(\mathbf{x}_k) - \mathbf{P}_{\mathbf{x}_{k-1}}^{\mathbf{x}_k} (\nabla f_{\mathcal{S}_2}(\mathbf{x}_{k-1}) - \mathbf{v}_{k-1}), \quad (3)$$

where the parallel transport $\mathbf{P}_{\mathbf{x}}^{\mathbf{z}}(\mathbf{y})$ (as defined in Section 2) transports \mathbf{y} from the tangent space at \mathbf{x} to that at the point \mathbf{z} . Here the parallel transport operation is necessary since $\nabla f_{\mathcal{S}_2}(\mathbf{x}_{k-1})$ and $\nabla f_{\mathcal{S}_2}(\mathbf{x}_k)$ are located in different tangent spaces. Given the gradient estimate \mathbf{v}_k , the variable is updated via normalized gradient descent $\mathbf{x}_{k+1} = \text{Exp}_{\mathbf{x}_k}(-\eta_k \frac{\mathbf{v}_k}{\|\mathbf{v}_k\|})$. Note that R-SRG [15] applies a similar recursion form as in (3) for full Riemannian gradient estimation, and the core difference between their method and ours lies in that R-SPIDER is equipped with gradient normalization which is missing in R-SRG. Then by carefully setting the learning rate η and mini-batch sizes of \mathcal{S}_1 and \mathcal{S}_2 , R-SPIDER only needs to sample a necessary number of data points for accurately estimating Riemannian gradient and sufficiently decreasing the objective at each iteration. In this way, R-SPIDER achieves sharper bounds of incremental first order oracle (IFO, see Definition 2) complexity than R-SRG and other Riemannian non-convex optimization methods.

Table 1 summarizes our main results on the computational complexity of R-SPIDER for non-convex problems, along with those for the aforementioned Riemannian algorithms. The following are some highlighted advantages of our results over the state-of-the-arts.

For the finite-sum setting of problem (1) with general non-convex functions, the IFO complexity of R-SPIDER to achieve $\mathbb{E}[\|\nabla f(\mathbf{x})\|] \leq \epsilon$ is $\mathcal{O}\left(\min\left(n + \frac{L\sqrt{n}}{\epsilon^2}, \frac{L\sigma}{\epsilon^3}\right)\right)$ which matches the lower IFO complexity bound in Euclidean space [30]. By comparison, the IFO complexity bounds of R-SRG and R-SVRG are $\mathcal{O}\left(n + \frac{L^2}{\epsilon^4}\right)$ and $\mathcal{O}\left(n + \frac{\zeta n^{\frac{2}{3}}}{\epsilon^{\frac{2}{3}}}\right)$, respectively. It can be verified that R-SPIDER improves over R-SRG by a factor of $\mathcal{O}\left(\frac{1}{\epsilon}\right)$ and R-SVRG by a factor $\mathcal{O}\left(n^{1/6}\right)$.

When $f(\mathbf{x})$ is a τ -gradient dominated function with finite-sum structure, R-SPIDER enjoys the IFO complexity of $\mathcal{O}(\min((n + \tau L\sqrt{n}) \log(\frac{1}{\epsilon}), \frac{\tau L\sigma}{\epsilon}))$ which is again lower than the $\mathcal{O}((n + \tau L\zeta^{\frac{1}{2}} n^{\frac{2}{3}}) \log(\frac{1}{\epsilon}))$ bound for R-SVRG by a factor of $\mathcal{O}(n^{1/6})$. Note that our IFO complexity is not dependent on the curvature parameter $\zeta(\geq 1)$ of the manifold \mathcal{M} , because our analysis does not involve the geodesic trigonometry inequality on a manifold. Compared with R-SRG with complexity bound $\mathcal{O}((n + \tau^2 L^2) \log(\frac{1}{\epsilon}))$, R-SPIDER is more efficient than R-SRG in large-sample-moderate-accuracy settings, e.g., in cases when n dominates $1/\epsilon$.

For the online version of problem (1), we establish the IFO complexity bounds $\mathcal{O}(\frac{L\sigma}{\epsilon^3})$ and $\mathcal{O}(\frac{\tau L\sigma}{\epsilon})$ for generic non-convex and gradient dominated problems, respectively. To our best knowledge, these non-asymptotic convergence results are novel to non-convex online Riemannian optimization. Comparatively, Bonnabel *et al.* [22] only provided asymptotic convergence analysis of R-SGD: the iterating sequence generated by R-SGD converges to a critical point when the iteration number approaches infinity.

Finally, our analysis reveals as a byproduct that R-SPIDER provably benefits from mini-batching. Specifically, our theoretic results imply linear speedups in parallel computing setting for large mini-batch sizes. We are not aware of any similar linear speedup results in the prior Riemannian stochastic algorithms.

Recently, the concurrent work [32] also applies SPIDER technique to solve problem (1) and achieves similar results. But this work differs from [32]. (1) Our algorithm uses an adaptive learning rate and a constant mini-batch size for general non-convex problems and employs a constant rate and an adaptive mini-batch size for gradient-dominated problems, while Zhang *et al.* [32] always adopt a constant step size and an adaptive mini-batch size. (2) For general non-convex problems, our computational complexity is $\mathcal{O}(\min(n + \frac{L\sqrt{n}}{\epsilon^2}, \frac{L\sigma}{\epsilon^3}))$ and is better than the complexity $\mathcal{O}(n + \frac{L\sqrt{n}}{\epsilon^2})$ in [32] when n dominates $\mathcal{O}(\frac{1}{\epsilon^2})$. For gradient-dominated problems, our work enjoys similar advantages. See details in Section 3.2.

2 Preliminaries

Throughout this paper, we assume that the Riemannian manifold $(\mathcal{M}, \mathbf{g})$ is a real smooth manifold \mathcal{M} equipped with a Riemannian metric \mathbf{g} . We denote the induced inner product $\langle \mathbf{y}, \mathbf{z} \rangle$ of any two vectors \mathbf{y} and \mathbf{z} in the tangent space $T_{\mathbf{x}}\mathcal{M}$ at the point \mathbf{x} as $\langle \mathbf{y}, \mathbf{z} \rangle = \mathbf{g}(\mathbf{y}, \mathbf{z})$, and denote the norm $\|\mathbf{y}\|$ as $\|\mathbf{y}\| = \sqrt{\mathbf{g}(\mathbf{y}, \mathbf{y})}$. Let $\nabla f_i(\mathbf{x})$ be the stochastic Riemannian gradient of $f_i(\mathbf{x})$ and also be a unbiased estimate to the full Riemannian

gradient $\nabla f(\mathbf{x})$, i.e. $\mathbb{E}_i[\nabla f_i(\mathbf{x})] = \nabla f(\mathbf{x})$.

The exponential mapping $\text{Exp}_{\mathbf{x}}(\mathbf{y})$ maps $\mathbf{y} \in T_{\mathbf{x}}\mathcal{M}$ to $\mathbf{z} \in \mathcal{M}$ such that there is a geodesic $\gamma(t)$ with $\gamma(0) = \mathbf{x}$, $\gamma(1) = \mathbf{z}$ and $\dot{\gamma}(0) = \frac{d}{dt}\gamma(t) = \mathbf{y}$. Here the geodesic $\gamma(t)$ is a constant speed curve $\gamma: [0, 1] \rightarrow \mathcal{M}$ which is locally distance minimized. If there exists a unique geodesic between any two points on \mathcal{M} , then the exponential map has an inverse mapping $\text{Exp}_{\mathbf{x}}^{-1}: \mathcal{M} \rightarrow T_{\mathbf{x}}\mathcal{M}$ and the geodesic is the unique shortest path with the geodesic distance $d(\mathbf{x}, \mathbf{z}) = \|\text{Exp}_{\mathbf{x}}^{-1}(\mathbf{z})\| = \|\text{Exp}_{\mathbf{z}}^{-1}(\mathbf{x})\|$ between $\mathbf{x}, \mathbf{z} \in \mathcal{M}$.

To utilize the historical and current Riemannian gradients, we need to transport the historical gradients into the tangent space of the current point such that these gradients can be linearly combined in one tangent space. For this purpose, we need to define the parallel transport operator $P_{\mathbf{x}}^{\mathbf{z}}: T_{\mathbf{x}}\mathcal{M} \rightarrow T_{\mathbf{z}}\mathcal{M}$ which maps $\mathbf{y} \in T_{\mathbf{x}}\mathcal{M}$ to $P_{\mathbf{x}}^{\mathbf{z}}(\mathbf{y}) \in T_{\mathbf{z}}\mathcal{M}$ while preserving the inner product and norm, i.e., $\langle \mathbf{y}_1, \mathbf{y}_2 \rangle = \langle P_{\mathbf{x}}^{\mathbf{z}}(\mathbf{y}_1), P_{\mathbf{x}}^{\mathbf{z}}(\mathbf{y}_2) \rangle$ and $\|\mathbf{y}\| = \|P_{\mathbf{x}}^{\mathbf{z}}(\mathbf{y})\|$ for $\forall \mathbf{y}_1, \mathbf{y}_2, \mathbf{y} \in T_{\mathbf{x}}\mathcal{M}$.

We impose on the loss components $f_i(\mathbf{x})$ the assumption of geodesic gradient-Lipschitz-smoothness. Such a smoothness condition is conventionally assumed in analyzing Riemannian gradient algorithms [15, 16, 33, 34].

Assumption 1 (Geodesically L -gradient-Lipschitz). *Each loss $f_i(\mathbf{x})$ is geodesically L -gradient Lipschitz such that $\mathbb{E}_i\|\nabla f_i(\mathbf{x}) - P_{\mathbf{y}}^{\mathbf{x}}(\nabla f_i(\mathbf{y}))\|^2 \leq L^2\|\text{Exp}_{\mathbf{x}}^{-1}(\mathbf{y})\|^2$.*

It can be shown that if each $f_i(\mathbf{x})$ is geodesically L -gradient-Lipschitz, then for any $\mathbf{x}, \mathbf{y} \in \mathcal{M}$,

$$f(\mathbf{y}) \leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \text{Exp}_{\mathbf{x}}^{-1}(\mathbf{y}) \rangle + \frac{L}{2}\|\text{Exp}_{\mathbf{x}}^{-1}(\mathbf{y})\|^2.$$

We also need to impose the following boundness assumption on the variance of stochastic gradient.

Assumption 2 (Bounded Stochastic Gradient Variance). *For any $\mathbf{x} \in \mathcal{M}$, the gradient variance of each loss $f_i(\mathbf{x})$ is bounded as $\mathbb{E}_i\|\nabla f_i(\mathbf{x}) - \nabla f(\mathbf{x})\|^2 \leq \sigma^2$.*

We further introduce τ -gradient dominated function [35, 36] which will also be investigated in this paper.

Definition 1 (τ -Gradient Dominated Functions). *$f(\mathbf{x})$ is said to be a τ -gradient dominated function if it satisfies $f(\mathbf{x}) - f(\mathbf{x}_*) \leq \tau\|\nabla f(\mathbf{x})\|^2$ for any $\mathbf{x} \in \mathcal{M}$, where τ is a universal constant and $\mathbf{x}_* = \text{argmin}_{\mathbf{x} \in \mathcal{M}} f(\mathbf{x})$ is the global minimizer of $f(\mathbf{x})$ on the manifold \mathcal{M} .*

The following defined incremental first order oracle (IFO) complexity is usually adopted as the computational complexity measurement for evaluating stochastic optimization algorithms [15, 16, 23, 24].

Definition 2 (IFO Complexity). *For $f(\mathbf{x})$ in problem (1), an IFO takes in an index $i \in [n]$ and a point \mathbf{x} , and returns the pair $(f_i(\mathbf{x}), \nabla f_i(\mathbf{x}))$.*

3 Riemannian SPIDER Algorithm

We first introduce the algorithm and then analyze its convergence behavior on general non-convex and gradient-dominated problems.

3.1 Algorithm

The R-SPIDER method is outlined in Algorithm 1. At its core, R-SPIDER customizes SPIDER to recursively estimate/track the full Riemannian gradient in a computationally economic way. For each cycle of p iterations, R-SPIDER first samples a large data batch \mathcal{S}_1 by with-replacement sampling and views the gradient estimate $\mathbf{v}_k = \nabla f_{\mathcal{S}_1}(\mathbf{x}_k) = \frac{1}{|\mathcal{S}_1|} \sum_{i \in \mathcal{S}_1} f_i(\mathbf{x}_k)$ as the snapshot gradient. For the next forthcoming $p-1$ iterations, R-SPIDER only samples a smaller mini-batch \mathcal{S}_2 and estimates the full Riemannian gradient $\nabla f(\mathbf{x}_k)$ as $\mathbf{v}_k = \nabla f_{\mathcal{S}_2}(\mathbf{x}_k) - \mathbf{P}_{\mathbf{x}_{k-1}}^{\mathbf{x}_k}(\nabla f_{\mathcal{S}_2}(\mathbf{x}_{k-1}) - \mathbf{v}_{k-1})$. Here the parallel transport operator $\mathbf{P}_{\mathbf{x}_{k-1}}^{\mathbf{x}_k}(\cdot)$ is applied to ensure that the Riemannian gradients can be linearly combined in a common tangent space. Then R-SPIDER performs normalized gradient descent to update $\mathbf{x}_{k+1} = \text{Exp}_{\mathbf{x}_k}(-\eta_k \frac{\mathbf{v}_k}{\|\mathbf{v}_k\|})$ until the termination of the algorithm.

The idea of recursive Riemannian gradient estimation has also been exploited by R-SRG [15]. Although sharing a similar spirit in full gradient approximation, R-SPIDER departs notably from R-SRG: at each iteration, R-SPIDER normalizes the gradient \mathbf{v}_k and thus is able to well control the distance $d(\mathbf{x}_k, \mathbf{x}_{k+1})$ between \mathbf{x}_k and \mathbf{x}_{k+1} by properly controlling the stepsize η , while R-SRG directly updates the variable without gradient normalization. It turns out that this normalization step is key to achieving faster convergence speed for non-convex problem in R-SPIDER, since it helps reduce the variance of stochastic gradient estimation by properly controlling the distance $d(\mathbf{x}_k, \mathbf{x}_{k+1})$ (see Lemma 1). As a consequence, at each iteration, R-SPIDER only needs to sample a necessary number of data points to estimate Riemannian gradient and decrease the objective sufficiently (see Theorems 1 and 2). In this way, R-SPIDER achieves lower overall computational complexity for solving problem (1).

3.2 Computational complexity analysis

The vanilla SPIDER is known to achieve nearly optimal iteration complexity bounds for stochastic non-convex optimization in Euclidean space [30]. We here show that R-SPIDER generalizes such an appealing property of SPIDER to Riemannian manifolds. We first present the following key lemma which guarantees sufficiently accurate Riemannian gradient estimation for R-SPIDER. We denote $\mathbb{I}_{\{\mathcal{E}\}}$ as the indicator function: if

Algorithm 1 R-SPIDER ($\mathbf{x}_0, \epsilon, \eta, p, |\mathcal{S}_1|, |\mathcal{S}_2|$)

- 1: **Input:** initialization \mathbf{x}_0 , accuracy ϵ , learning rate η , iteration interval p , mini-batch sizes $|\mathcal{S}_1|$ and $|\mathcal{S}_2|$.
 - 2: **for** $k = 0$ to $K - 1$ **do**
 - 3: **if** $\text{mod}(k, p) = 0$ **then**
 - 4: Draw mini-batch \mathcal{S}_1 and compute $\mathbf{v}_k = \nabla f_{\mathcal{S}_1}(\mathbf{x}_k)$;
 - 5: **else**
 - 6: Draw mini-batch \mathcal{S}_2 and compute $\nabla f_{\mathcal{S}_2}(\mathbf{x}_k)$;
 - 7: $\mathbf{v}_k = \nabla f_{\mathcal{S}_2}(\mathbf{x}_k) - \mathbf{P}_{\mathbf{x}_{k-1}}^{\mathbf{x}_k}(\nabla f_{\mathcal{S}_2}(\mathbf{x}_{k-1}) - \mathbf{v}_{k-1})$;
 - 8: **end if**
 - 9: $\mathbf{x}_{k+1} = \text{Exp}_{\mathbf{x}_k}(-\eta_k \frac{\mathbf{v}_k}{\|\mathbf{v}_k\|})$;
 - 10: **end for**
 - 11: **Output:** $\tilde{\mathbf{x}}$ which is chosen uniformly at random from $\{\mathbf{x}_k\}_{k=0}^{K-1}$.
-

the event \mathcal{E} is true, then $\mathbb{I}_{\{\mathcal{E}\}} = 1$; otherwise, $\mathbb{I}_{\{\mathcal{E}\}} = 0$.

Lemma 1 (Bounded Gradient Estimation Error). *Suppose Assumptions 1 and 2 hold. Let $k_0 = \lfloor k/p \rfloor$ and $\tilde{k}_0 = k_0 p$. The estimation error between the full Riemannian gradient $\nabla f(\mathbf{x}_k)$ and its estimate \mathbf{v}_k in Algorithm 1 is bounded as*

$$\begin{aligned} & \mathbb{E} \left[\|\mathbf{v}_k - \nabla f(\mathbf{x}_k)\|^2 \mid \mathbf{x}_{\tilde{k}_0}, \dots, \mathbf{x}_{\tilde{k}_0+p-1} \right] \\ & \leq \mathbb{I}_{\{|\mathcal{S}_1| < n\}} \frac{\sigma^2}{|\mathcal{S}_1|} + \frac{L^2}{|\mathcal{S}_2|} \sum_{i=\tilde{k}_0}^{\tilde{k}_0+p-1} d^2(\mathbf{x}_i, \mathbf{x}_{i+1}), \end{aligned}$$

where $d(\mathbf{x}_i, \mathbf{x}_{i+1})$ is the distance between \mathbf{x}_i and \mathbf{x}_{i+1} .

Proof. The key is to carefully handle the exponential mapping and parallel transport operators introduced for vector computation. See details in Appendix A. \square

Lemma 1 tells that by properly selecting the mini-batch sizes $|\mathcal{S}_1|$ and $|\mathcal{S}_2|$, the accuracy of gradient estimate \mathbf{v}_k can be controlled. Benefiting from the normalization step, we have $d(\mathbf{x}_k, \mathbf{x}_{k+1}) = \|\text{Exp}_{\mathbf{x}_k}^{-1}(\mathbf{x}_{k+1})\| = \eta_k$. As a result, the gradient estimation error can be bounded as $\mathbb{E}[\|\mathbf{v}_k - \nabla f(\mathbf{x}_k)\|^2 \mid \mathbf{x}_{\tilde{k}_0}, \dots, \mathbf{x}_{\tilde{k}_0+p-1}] \leq \mathbb{I}_{\{|\mathcal{S}_1| < n\}} \frac{\sigma^2}{|\mathcal{S}_1|} + \frac{L^2}{|\mathcal{S}_2|} \sum_{i=\tilde{k}_0}^{\tilde{k}_0+p-1} \eta_i^2$. Then we are able to analyze the rate-of-convergence of R-SPIDER.

Finite-sum setting. We first consider problem (1) under finite-sum setting. By properly selecting parameters, we prove that at each iteration, the sequence $\{\mathbf{x}_k\}$ produced by Algorithm 1 can lead to sufficient decrease of the objective loss $f(\mathbf{x})$ when $\|\mathbf{v}_k\|$ is large. Based on this results, we further derive the iteration number of Algorithm 1 for computing an ϵ -accuracy solution. The result is formally summarized in Theorem 1.

Theorem 1. *Suppose Assumptions 1 and 2 hold. Let $s = \min(n, \frac{16\sigma^2}{\epsilon^2})$, $p = n_0 s^{\frac{1}{2}}$, $\eta_k = \min(\frac{\epsilon}{2Ln_0}, \frac{\|\mathbf{v}_k\|}{4Ln_0})$,*

$|\mathcal{S}_1| = s$, $|\mathcal{S}_2| = \frac{4s^{\frac{1}{2}}}{n_0}$ and $n_0 \in [1, 4s^{\frac{1}{2}}]$. Then for finite-sum problem (1), the sequence $\{\mathbf{x}_k\}$ produced by Algorithm 1 satisfies

$$\mathbb{E}[f(\mathbf{x}_{k+1}) - f(\mathbf{x}_k)] \leq -\frac{\epsilon}{64Ln_0} (12\mathbb{E}[\|\mathbf{v}_k\|] - 7\epsilon).$$

Moreover, to achieve $\mathbb{E}[\|\nabla f(\tilde{\mathbf{x}})\|] \leq \epsilon$, Algorithm 1 will terminate at most $(\frac{14Ln_0\Delta}{\epsilon^2})$ iterations in expectation, where $\Delta = f(\mathbf{x}_0) - f(\mathbf{x}_*)$ with $\mathbf{x}_* = \operatorname{argmin}_{\mathbf{x} \in \mathcal{M}} f(\mathbf{x})$.

Proof. The result comes readily from geodesically L -gradient-Lipschitz of f and the variance bound in Lemma 1. See Appendix B.1 for a complete proof. \square

Theorem 1 shows that Algorithm 1 only needs to run at most $(\frac{14Ln_0\Delta}{\epsilon^2})$ iteration to compute an ϵ -accuracy solution $\tilde{\mathbf{x}}$, i.e. $\mathbb{E}[\|\nabla f(\tilde{\mathbf{x}})\|] \leq \epsilon$. This means the convergence rate of R-SPIDER is at the order of $\mathcal{O}(\frac{Ln_0\Delta}{\epsilon^2})$. Besides, by one iteration loop of Algorithm 1, the objective value $f(\mathbf{x}_k)$ monotonously decreases in expectation when $\mathbb{E}[\|\mathbf{v}_k\|]$ is large, e.g. $\mathbb{E}[\|\mathbf{v}_k\|] \geq \frac{7\epsilon}{12}$. By comparison, Kasai *et al.* [15] only proved the sublinear convergence rate of the gradient norm $\mathbb{E}[\|\nabla f(\mathbf{x})\|^2]$ in R-SRG and did not reveal any convergence behavior of the objective $f(\mathbf{x})$. Moreover, Theorem 1 yields as a byproduct the benefits of mini-batching to R-SPIDER. Indeed, by controlling the parameter n_0 in R-SPIDER, the mini-batch size $|\mathcal{S}_2|$ at each iteration can range from 1 to $\min(4\sqrt{n}, \frac{16\sigma}{\epsilon})$. Also, it can be seen from Theorem 1 that larger mini-batch size allows more aggressive step size η_k and thus leads to less necessary iterations to achieve an ϵ -accuracy solution. More specifically, the convergence rate bound $\mathcal{O}(\frac{Ln_0\Delta}{\epsilon^2})$ indicates that at least in theory, increasing the mini-batch sizes in R-SPIDER provides linear speedups in parallel computing environment. In contrast, these important benefits of mini-batching are not explicitly analyzed in the existing Riemannian stochastic gradient algorithms [15, 16].

Based on Theorem 1, we can derive the IFO complexity of R-SPIDER for non-convex problems in Corollary 1.

Corollary 1. *Using the same assumptions and parameters in Theorem 1, the IFO complexity of Algorithm 1 is $\mathcal{O}(\min(n + \frac{L\Delta\sqrt{n}}{\epsilon^2}, \frac{L\Delta\sigma}{\epsilon^3}))$ for achieving $\mathbb{E}[\|\nabla f(\tilde{\mathbf{x}})\|] \leq \epsilon$.*

Proof. The result is obtained directly from a cumulation of IFOs at each step of iteration. See Appendix B.2. \square

From Corollary 1, the IFO complexity of R-SPIDER for non-convex finite-sum problems is at the order of $\mathcal{O}(\frac{1}{\epsilon^2} \min(\sqrt{n}, \frac{1}{\epsilon}))$. This result matches the state-of-the-art complexity bounds for general non-convex optimization problems in Euclidean space [30, 37]. Indeed,

under Assumption 1, Fang *et al.* [30] proved that the lower IFO complexity bound for finite-sum problem (1) in Euclidean space is $\mathcal{O}(n + \frac{L\Delta\sqrt{n}}{\epsilon^2})$ when the number n of the component function obeys $n \leq \mathcal{O}(\frac{L^2\Delta^2}{\epsilon^4})$. In the sense that Euclidean space is a special case of Riemannian manifold, our IFO complexity $\mathcal{O}(n + \frac{L\Delta\sqrt{n}}{\epsilon^2})$ for finite-sum problem (1) under Assumption 1 is nearly optimal. If Assumption 2 holds in addition, we can establish tighter IFO complexity $\mathcal{O}(\frac{1}{\epsilon^2} \min(\sqrt{n}, \frac{1}{\epsilon}))$. This is because when the sample number n satisfies $n \geq \frac{16\sigma^2}{\epsilon^2}$, by sampling $|\mathcal{S}_1| = \frac{16\sigma^2}{\epsilon^2}$ and $|\mathcal{S}_2| = \frac{16\sigma}{n_0\epsilon}$, the gradient estimation error already satisfies $\mathbb{E}[\|\mathbf{v}_k - \nabla f(\mathbf{x}_k)\|^2] \leq \frac{\epsilon^2}{8}$. Accordingly, if $\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}\|\mathbf{v}_k\| \leq 0.5\epsilon$ which is actually achieved after K iterations, then $\mathbb{E}[\|\nabla f(\tilde{\mathbf{x}})\|] = \frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}\|\nabla f(\mathbf{x}_k)\| \leq \frac{1}{K} \sum_{k=0}^{K-1} (\mathbb{E}\|\nabla f(\mathbf{x}_k) - \mathbf{v}_k\| + \mathbb{E}\|\mathbf{v}_k\|) \leq \epsilon$. So here it is only necessary to sample $|\mathcal{S}_1| = \frac{16\sigma^2}{\epsilon^2}$ data points instead of the entire set of n samples.

Kasai *et al.* [15] proved that the IFO complexity of R-SRG is at the order of $\mathcal{O}(n + \frac{L^2}{\epsilon^4})$ to obtain an ϵ -accuracy solution. By comparison, we prove that R-SPIDER enjoys the complexity of $\mathcal{O}(\frac{1}{\epsilon^2} \min(\sqrt{n}, \frac{1}{\epsilon}))$, which is at least lower than R-SRG by a factor of $\frac{1}{\epsilon}$. This is because the normalization step in R-SPIDER allows us to well control the gradient estimation error and thus avoids sampling too many redundant samples at each iteration, resulting in sharper IFO complexity. Zhang *et al.* [16] showed that R-SVRG has the IFO complexity $\mathcal{O}(n + \frac{\zeta^{1/2}n^{2/3}}{\epsilon^2})$, where $\zeta \geq 1$ denotes the curvature parameter. Therefore, R-SPIDER improves over R-SVRG by a factor at least $n^{1/6}$ in IFO complexity. Note, here the curvature parameter ζ does not appear in our bounds, since we have avoided using the trigonometry inequality which characterizes the trigonometric geometric in Riemannian manifold [16, 21, 22].

Online setting. Next we consider the online setting of problem (1). Similar to finite-sum setting, we prove in Theorem 2 that the objective $f(\mathbf{x})$ can be sufficiently decreased when the gradient norm is not too small.

Theorem 2. *Suppose Assumptions 1 and 2 hold. Let $p = \frac{\sigma n_0}{\epsilon}$, $\eta_k = \min(\frac{\epsilon}{2Ln_0}, \frac{\|\mathbf{v}_k\|}{4Ln_0})$, $|\mathcal{S}_1| = \frac{64\sigma^2}{\epsilon^2}$, $|\mathcal{S}_2| = \frac{4\sigma}{\epsilon n_0}$ and $n_0 \in [1, \frac{4\sigma}{\epsilon}]$. For problem (1) under online setting, the sequence $\{\mathbf{x}_k\}$ produced by Algorithm 1 satisfies*

$$\mathbb{E}[f(\mathbf{x}_{k+1}) - f(\mathbf{x}_k)] \leq -\frac{\epsilon}{64Ln_0} (12\mathbb{E}[\|\mathbf{v}_k\|] - 7\epsilon).$$

Moreover, to achieve $\mathbb{E}[\|\nabla f(\tilde{\mathbf{x}})\|] \leq \epsilon$, Algorithm 1 will terminate at most $(\frac{14Ln_0\Delta}{\epsilon^2})$ iterations in expectation, where $\Delta = f(\mathbf{x}_0) - f(\mathbf{x}_*)$ with $\mathbf{x}_* = \operatorname{argmin}_{\mathbf{x} \in \mathcal{M}} f(\mathbf{x})$.

Proof. The proof mimics that of Theorem 1 with proper adaptation to online setting. See Appendix B.3. \square

As a direct consequence of this result, the following corollary establishes the IFO complexity of R-SPIDER for the online optimization.

Corollary 2. *Using the same assumptions and parameters in Theorem 2, the IFO complexity of Algorithm 1 is $\mathcal{O}\left(\frac{L\sigma\Delta}{\epsilon^3}\right)$ to achieve $\mathbb{E}[\|\nabla f(\mathbf{x})\|] \leq \epsilon$.*

Proof. See Appendix B.4 for a proof of this result. \square

Bonnabel *et al.* [22] have also analyzed R-SGD under online setting, but only with asymptotic convergence guarantee obtained. By comparison, we for the first time establish non-asymptotic complexity bounds for Riemannian online non-convex optimization.

Algorithm 2 Riemannian Gradient Dominated SPIDER (R-GD-SPIDER)

- 1: **Input:** initial point $\tilde{\mathbf{x}}_0$, initial accuracy ϵ_0 , learning rate η^0 , mini-batch sizes $|\mathcal{S}_1^0|$ and $|\mathcal{S}_2^0|$, iteration interval p^0 , final accuracy ϵ
 - 2: **for** $t = 1$ to T **do**
 - 3: $\tilde{\mathbf{x}}_t = \text{R-SPIDER}(\tilde{\mathbf{x}}_{t-1}, \epsilon_{t-1}, \eta^t, p^t, |\mathcal{S}_1^t|, |\mathcal{S}_2^t|)$.
 - 4: Set $\epsilon_t = 0.5\epsilon_{t-1}$, and $\eta^t, p^t, |\mathcal{S}_1^t|, |\mathcal{S}_2^t|$ properly.
 - 5: **end for**
 - 6: **Output:** $\tilde{\mathbf{x}}_t$
-

3.3 On gradient dominated functions

We now turn to a special case of problem (1) with gradient dominated loss function as defined in Definition 1. For instance, the strongly geodesically convex (SGC) functions¹ are gradient dominated. Some non-strongly convex problems, *e.g.* ill-conditioned linear prediction and logistic regression [38], and Riemannian non-convex problems, *e.g.* PCA [16], also belong to gradient dominated functions. Please refer to [36, 38] for more instances of gradient dominated functions. To better fit gradient dominated functions, we develop the Riemannian gradient dominated SPIDER (R-GD-SPIDER) as a multi-stage variant of R-SPIDER. A high-level description of R-GD-SPIDER is outlined in Algorithm 2. The basic idea is to use more aggressive learning rates in early stage of processing and gradually shrink the learning rate in later stage. With the help of such a simulated annealing process, R-GD-SPIDER exhibits linear convergence behavior for finite-sum problems, as formally stated in Theorem 3. For the t -th iteration in Algorithm 2, R-SPIDER uses $|\mathcal{S}_1^t|$ and $|\mathcal{S}_{2,k}^t|$ samples to compute \mathbf{v}_k for $k = \lfloor k/p \rfloor \cdot p$ and $k \neq \lfloor k/p \rfloor \cdot p$, respectively.

¹A strongly geodesically convex function satisfies $f(\mathbf{y}) \geq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \text{Exp}_{\mathbf{x}}^{-1}(\mathbf{y}) \rangle + \frac{\mu}{2} \|\text{Exp}_{\mathbf{x}}^{-1}(\mathbf{y})\|^2, \forall \mathbf{x}, \mathbf{y} \in \mathcal{M}$, for some $\mu > 0$, which immediately implies $f(\mathbf{x}) - f(\mathbf{x}_*) \leq \frac{1}{2\mu} \|\nabla f(\mathbf{x})\|^2$ by Cauchy-Schwarz inequality.

Theorem 3. *Suppose that function $f(\mathbf{x})$ is τ -gradient dominated, and Assumptions 1 and 2 hold. For finite-sum setting, at the t -th iteration, set $\epsilon_0 = \frac{\sqrt{\Delta}}{2\sqrt{\tau}}$, $\epsilon_t = \frac{\epsilon_0}{2^t}$, $s_t = \min(n, \frac{32\sigma^2}{\epsilon_{t-1}^2})$, $p^t = n_0^t \sqrt{s_t}$, $\eta_k^t = \frac{\|\mathbf{v}_k^t\|}{2Ln_0^t}$, $|\mathcal{S}_1^t| = s_t$, $|\mathcal{S}_{2,k}^t| = \min(\frac{8p^t \|\mathbf{v}_{k-1}^t\|^2}{(n_0^t)^2 \epsilon_{t-1}^2}, n)$, where $n_0^t \in [1, \frac{8\sqrt{s_t} \|\mathbf{v}_{k-1}^t\|^2}{\epsilon_{t-1}^2}]$. (1) The sequence $\{\tilde{\mathbf{x}}_t\}$ produced by Algorithm 2 satisfies*

$$\mathbb{E}[f(\tilde{\mathbf{x}}_t) - f(\mathbf{x}_*)] \leq \frac{\Delta}{4^t} \quad \text{and} \quad \mathbb{E}[\|\nabla f(\tilde{\mathbf{x}}_t)\|] \leq \frac{1}{2^t} \sqrt{\frac{\Delta}{\tau}},$$

where $\Delta = f(\tilde{\mathbf{x}}_0) - f(\mathbf{x}_*)$ with $\mathbf{x}_* = \text{argmin}_{\mathbf{x} \in \mathcal{M}} f(\mathbf{x})$. (2) To achieve $\mathbb{E}[\|\nabla f(\tilde{\mathbf{x}}_T)\|] \leq \epsilon$, in expectation the IFO complexity is $\mathcal{O}\left(\min\left((n + \tau L\sqrt{n}) \log\left(\frac{1}{\epsilon}\right), \frac{\tau L\sigma}{\epsilon}\right)\right)$.

Proof. The part (1) follows immediately from the update rule of ϵ_t . The part (2) can be proved by establishing the IFO bound $\min\left(n + \tau L\sqrt{n}, \frac{\tau L\sigma}{\epsilon_{t+1}}\right)$ for each stage t and then putting them together. See Appendix C.1. \square

The main message conveyed by Theorem 3 is that R-GD-SPIDER enjoys a global linear rate of convergence and its IFO complexity is at the order of $\mathcal{O}\left(\min\left((n + \tau L\sqrt{n}) \log\left(\frac{1}{\epsilon}\right), \frac{\tau L\sigma}{\epsilon}\right)\right)$. For R-SVRG with τ -gradient dominated functions, Zhang *et al.* [16] also established a linear convergence rate and an IFO complexity bound $\mathcal{O}\left((n + \tau L\zeta^{\frac{1}{2}} n^{\frac{3}{2}}) \log\left(\frac{1}{\epsilon}\right)\right)$. As a comparison, our R-GD-SPIDER makes an improvement over R-SVRG in IFO complexity by a factor of $n^{\frac{1}{6}}$. For R-SRG [15], the corresponding IFO complexity is $\mathcal{O}\left((n + \tau^2 L^2) \log\left(\frac{1}{\epsilon}\right)\right)$. Therefore, in terms of IFO complexity, R-GD-SPIDER is superior to R-SRG when the optimization accuracy ϵ is moderately small at a huge data size n .

Turning to the online setting, R-GD-SPIDER also converges linearly, as formally stated in Theorem 4.

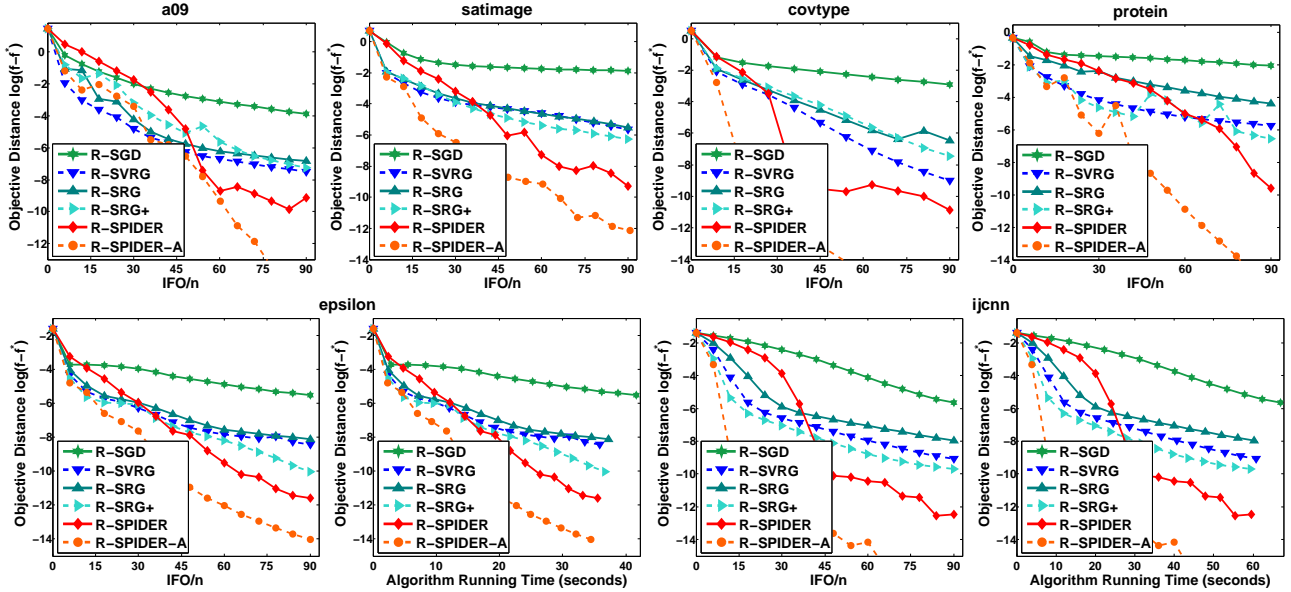
Theorem 4. *Suppose that $f(\mathbf{x})$ is τ -gradient dominated, and Assumptions 1 and 2 hold. For online setting, at the t -th iteration, let $\epsilon_0 = \frac{\sqrt{\Delta}}{2\sqrt{\tau}}$, $\epsilon_t = \frac{\epsilon_0}{2^t}$, $p_t = \frac{\sigma n_0^t}{\epsilon_{t-1}}$, $\eta_k^t = \frac{\|\mathbf{v}_k^t\|}{2Ln_0^t}$, $|\mathcal{S}_1^t| = \frac{32\sigma^2}{\epsilon_{t-1}^2}$, $|\mathcal{S}_{2,k}^t| = \frac{8\sigma \|\mathbf{v}_{k-1}^t\|^2}{\epsilon_{t-1}^3 n_0^t}$, where $n_0^t \in [1, \frac{8\sigma \|\mathbf{v}_{k-1}^t\|^2}{\epsilon_{t-1}^3}]$.*

(1) The sequence $\{\tilde{\mathbf{x}}_t\}$ produced by Algorithm 2 satisfies

$$\mathbb{E}[f(\tilde{\mathbf{x}}_t) - f(\mathbf{x}_*)] \leq \frac{\Delta}{4^t} \quad \text{and} \quad \mathbb{E}[\|\nabla f(\tilde{\mathbf{x}}_t)\|] \leq \frac{1}{2^t} \sqrt{\frac{\Delta}{\tau}},$$

where $\Delta = f(\tilde{\mathbf{x}}_0) - f(\mathbf{x}_*)$ with $\mathbf{x}_* = \text{argmin}_{\mathbf{x} \in \mathcal{M}} f(\mathbf{x})$. (2) To achieve $\mathbb{E}[\|\nabla f(\tilde{\mathbf{x}}_T)\|] \leq \epsilon$, in expectation the IFO complexity is $\mathcal{O}\left(\frac{\tau L\sigma}{\epsilon}\right)$.

Proof. See Appendix C.2 for a proof of this result. \square

Figure 1: Comparison among Riemannian stochastic gradient algorithms on k -PCA problem.

Such a non-asymptotic convergence result is new to online Riemannian gradient dominated optimization.

4 Experiments

Here we compare the proposed R-SPIDER with several Riemannian stochastic gradient algorithms, including R-SGD [22], R-SVRG [16, 23], R-SRG [15] and R-SRG+ [15]. We evaluate all the considered algorithms on two learning tasks: the k -PCA problem and the low-rank matrix completion problem. We run simulations on ten datasets, including six datasets from LibSVM, three face datasets (YaleB, AR and PIE) and one recommendation dataset (MovieLens-1M). The details of these datasets are described in Appendix D.

A practical implementation of R-SPIDER. To achieve the IFO complexity in Corollary 1, it is suggested to set the learning rate as $\eta = \frac{\epsilon}{4Ln_0}$ where ϵ is the desired optimization accuracy. However, since in the initial epochs the computed point is far from the optimum to problem (1), using a tiny learning rate could usually be conservative. In contrast, by using a more aggressive learning rate at the initial optimization stage, we can expect stable but faster convergence behavior. Here for R-SPIDER we design a decaying learning rate with formulation $\eta_k = \alpha^{\lfloor \frac{k}{\beta} \rfloor} \cdot \beta$ and call it ‘‘R-SPIDER-A’’, where α and β are two constants. In our experiments, α is selected from $\{0.8, 0.85, 0.9, 0.95, 0.99\}$ and β from $\{5 \times 10^{-2}, 10^{-2}, 5 \times 10^{-3}, 10^{-3}\}$.

Evaluation on the k -PCA problem. Given n data points, k -PCA aims at computing their first k leading eigenvectors, which is formulated as $\min_{U \in \text{Gr}(k,d)} \frac{1}{n} \sum_{i=1}^n \mathbf{a}_i^\top U U^\top \mathbf{a}_i$, where $\mathbf{a}_i \in \mathbb{R}^d$ de-

notes the i -th sample vector and $\text{Gr}(k, d) = \{U \in \mathbb{R}^{d \times k} \mid U^\top U = I\}$ denotes the Grassmann manifold. For this problem, we can directly obtain the ground truth U^* by using singular value decomposition (SVD), and then use $f(U^*)$ as optimal value f^* for sub-optimality estimation in Figures 1 and 2. In this experiment, we compute the first ten leading eigenvectors.

From the experimental results in Figure 1, one can observe that our R-SPIDER-A converges significantly faster than other algorithms and R-SPIDER can also quickly converge to a relatively high accuracy, e.g. 10^{-8} . In the initial epochs, R-SPIDER is comparable to other algorithms, showing relatively flat convergence behavior, mainly due to its very small learning rate and gradient normalization. Then along with more iterations, the computed solution becomes close to the optimum. Accordingly, the gradient begins to vanish and those considered algorithms without normalization tend to update the variable with small progress. In contrast, thanks to the normalization step, R-SPIDER moves more rapidly along the gradient descent direction and thus has sharper convergence curves. Meanwhile, R-SPIDER-A uses a relatively more aggressive learning rate in the initial epochs and decreases the learning rate along with more iterations. As a result, it exhibits the sharpest convergence behavior. On epsilon and ijcn datasets (the bottom of Figure 1) we further plot the sub-optimality versus running-time curves. The main observations from this group of curves are consistent with those of IFO complexity, implying that the IFO complexity can comprehensively reflect the overall computational performance of a first-order Riemannian algorithm. See Figure 5 in Appendix D.2 for more experimental results on running time comparison.

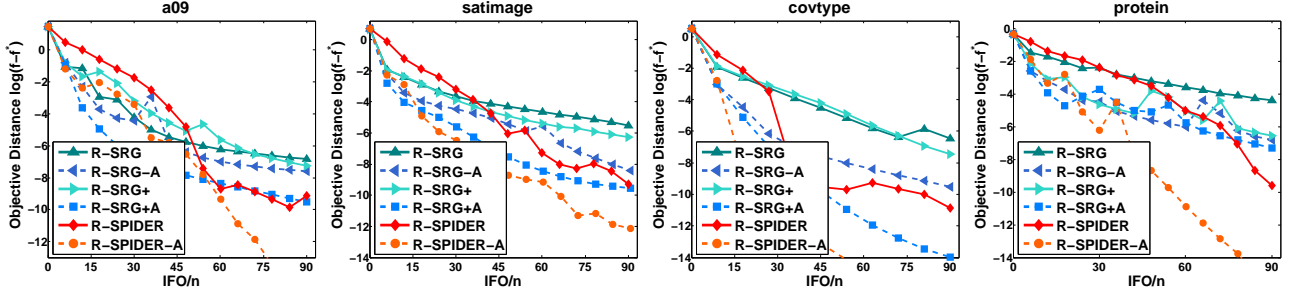


 Figure 2: Comparison between R-SPIDER and R-SRG with adaptive learning rates on k -PCA problem.


Figure 3: Comparison among Riemannian stochastic gradient algorithms on low-rank matrix completion problem.

In Figure 2, we compare R-SPIDER-A more closely with R-SRG-A and R-SRG+A which are respectively the counterparts of R-SRG and R-SRG+ with adaptive learning rate $\eta_k = \alpha(1 + \alpha\lambda_\alpha \lfloor \frac{k}{p} \rfloor)$ [15]. Here α and λ_α are tunable hyper-parameters. From the results, one can observe that the algorithm with adaptive learning rate usually outperforms the vanilla counterpart, which demonstrates the effectiveness of such an implementation trick. Moreover, R-SPIDER-A is consistently superior to R-SRG-A and R-SRG+A. See Figure 6 in Appendix D.3 for more results in this line.

Evaluation on the low-rank matrix completion (LRMC) problem. Given a low-rank incomplete observation $\mathbf{A} \in \mathbb{R}^{d \times n}$, the LRMC problem aims at exactly recovering \mathbf{A} . The mathematical formulation is $\min_{\mathbf{U} \in \text{Gr}(k,d), \mathbf{G}_i \in \mathbb{R}^{k \times n}} \|\mathcal{P}_\Omega(\mathbf{A}) - \mathcal{P}_\Omega(\mathbf{UG})\|^2$, where the set Ω of locations corresponds to the observed entries, namely $(i, j) \in \Omega$ if \mathbf{A}_{ij} is observed. \mathcal{P}_Ω is a linear operator that extracts entries in Ω and fills the entries not in Ω with zeros. The LRMC problem can be expressed equivalently as $\min_{\mathbf{U} \in \text{Gr}(k,d), \mathbf{G}_i \in \mathbb{R}^{k \times n}} \frac{1}{n} \sum_{i=1}^n \|\mathcal{P}_{\Omega_i}(\mathbf{A}_i) - \mathcal{P}_{\Omega_i}(\mathbf{UG}_i)\|^2$. Since there is no ground truth for the optimum, we run Riemannian GD sufficiently long until the gradient satisfies $\|\nabla f(\mathbf{x})\|/\|\mathbf{x}\| \leq 10^{-8}$, and then use the output as an approximate optimal value f^* for sub-optimality estimation in Figure 3. We test the considered algorithms on YaleB, AR, PIE and MovieLens-1M, considering these data approximately lie on a union of low-rank subspaces [15, 39]. For face images, we randomly sample 50% pixels in each image as the observations and set $k = 30$. For MovieLens-1M, we use its one million ratings for 3,952 movies from 6,040 users as the observations and set $k = 100$.

From Figure 3, R-SPIDER-A and R-SPIDER show very similar behaviors as those in Figure 1. More specifically, R-SPIDER-A achieves fastest convergence rate, and R-SPIDER has similar convergence speed as other algorithms in the initial epochs and then runs faster along with more epochs. All these results confirm the superiority of R-SPIDER and R-SPIDER-A.

5 Conclusions

We proposed R-SPIDER, which is an efficient Riemannian gradient method for non-convex stochastic optimization on Riemannian manifolds. Compared to existing first-order Riemannian algorithms, R-SPIDER enjoys provably lower computational complexity bounds for finite-sum minimization. For online optimization, similar non-asymptotic bounds are established for R-SPIDER, which to our best knowledge has not been addressed in previous study. For the special case of gradient dominated functions, we further developed a variant of R-SPIDER with improved linear rate of convergence. Numerical results confirm the computational superiority of R-SPIDER over the state-of-the-arts.

Acknowledgements

Jiashi Feng was partially supported by NUS startup R-263-000-C08-133, MOE Tier-I R-263-000-C21-112, NUS IDS R-263-000-C67-646, ECRA R-263-000-C87-133 and MOE Tier-II R-263-000-D17-112. Xiao-Tong Yuan is supported by Natural Science Foundation of China (NSFC) under Grant 61522308 and Grant 61876090.

References

- [1] S. Wold, K. Esbensen, and P. Geladi. Principal component analysis. *Chemometrics and intelligent laboratory systems*, 2(1-3):37–52, 1987. [1](#)
- [2] M. Tan, I. Tsang, L. Wang, B. Vandereycken, and S. Pan. Riemannian pursuit for big matrix recovery. In *Proc. Int'l Conf. Machine Learning*, pages 1539–1547, 2014. [1](#)
- [3] B. Vandereycken. Low-rank matrix completion by Riemannian optimization. *SIAM Journal on Optimization*, 23(2):1214–1236, 2013. [1](#)
- [4] P. Zhou and J. Feng. Outlier-robust tensor PCA. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 1–9, 2017. [1](#)
- [5] B. Mishra and R. Sepulchre. R3MC: A Riemannian three-factor algorithm for low-rank matrix completion. In *Proc. IEEE Conf. on Decision and Control*, pages 1137–1142, 2014. [1](#)
- [6] H. Kasai and B. Mishra. Low-rank tensor completion: a Riemannian manifold preconditioning approach. In *Proc. Int'l Conf. Machine Learning*, pages 1012–1021, 2016. [1](#)
- [7] P. Zhou, C. Lu, Z. Lin, and C. Zhang. Tensor factorization for low-rank tensor completion. *IEEE Transactions on Image Processing*, 27(3):1152 – 1163, 2017. [1](#)
- [8] P. Zhou, X. Yuan, and J. Feng. Efficient stochastic gradient hard thresholding. In *Proc. Conf. Neural Information Processing Systems*, 2018. [1](#)
- [9] A. Cherian and S. Sra. Riemannian dictionary learning and sparse coding for positive definite matrices. *IEEE trans. on Neural Networks and Learning Systems*, 28(12):2859–2871, 2017. [1](#)
- [10] P. Zhou, C. Fang, Z. Lin, C. Zhang, and E. Chang. Dictionary learning with structured noise. *Neurocomputing*, 273:414–423, 2017. [1](#)
- [11] J. Sun, Q. Qu, and J. Wright. Complete dictionary recovery over the sphere ii: Recovery by Riemannian trust-region method. *IEEE Trans. on Information Theory*, 63(2):885–914, 2017. [1](#)
- [12] P. Zhou, C. Zhang, and Z. Lin. Bilevel model based discriminative dictionary learning for recognition. *IEEE Trans. on Image Processing*, 26(3):1173–1187, 2017. [1](#)
- [13] R. Hosseini and S. Sra. Matrix manifold optimization for Gaussian mixtures. In *Proc. Conf. Neural Information Processing Systems*, pages 910–918, 2015. [1](#)
- [14] G. Meyer, S. Bonnabel, and R. Sepulchre. Linear regression under fixed-rank constraints: a Riemannian approach. In *Proc. Int'l Conf. Machine Learning*, 2011. [1](#)
- [15] H. Kasai, H. Sato, and B. Mishra. Riemannian stochastic recursive gradient algorithm with retraction and vector transport and its convergence analysis. In *Proc. Int'l Conf. Machine Learning*, pages 2521–2529, 2018. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [16] H. Zhang, S. Reddi, and S. Sra. Riemannian SVRG: Fast stochastic optimization on Riemannian manifolds. In *Proc. Conf. Neural Information Processing Systems*, pages 4592–4600, 2016. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [17] E. Oja. Principal components, minor components, and linear neural networks. *Neural networks*, 5(6):927–935, 1992. [1](#)
- [18] J. da Cruz Neto, L. De Lima, and P. Oliveira. Geodesic algorithms in Riemannian geometry. *Balkan Journal of Geometry and its Applications*, 3(2):89–100, 1998. [1](#)
- [19] R. Badeau, B. David, and G. Richard. Fast approximated power iteration subspace tracking. *IEEE Trans. on Signal Processing*, 53(8):2931–2941, 2005. [1](#)
- [20] H. Zhang and S. Sra. An estimate sequence for geodesically convex optimization. In *Proc. Conf. on Learning Theory*, pages 1703–1723, 2018. [1](#)
- [21] H. Zhang and S. Sra. First-order methods for geodesically convex optimization. In *Proc. Conf. on Learning Theory*, pages 1617–1638, 2016. [1](#), [5](#)
- [22] S. Bonnabel. Stochastic gradient descent on Riemannian manifolds. *IEEE Trans. Automatic Control*, 58(9):2217–2229, 2013. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [23] H. Kasai, H. Sato, and B. Mishra. Riemannian stochastic variance reduced gradient on Grassmann manifold. *arXiv preprint arXiv:1605.07367*, 2016. [1](#), [2](#), [3](#), [7](#)
- [24] H. Kasai, H. Sato, and B. Mishra. Riemannian stochastic quasi-Newton algorithm with variance reduction and its convergence analysis. *Prof. Int'l Conf. Artificial Intelligence and Statistics*, 2018. [1](#), [2](#), [3](#)
- [25] Y. Liu, F. Shang, J. Cheng, H. Cheng, and L. Jiao. Accelerated first-order methods for geodesically convex optimization on Riemannian manifolds. In *Proc. Conf. Neural Information Processing Systems*, pages 4868–4877, 2017. [1](#)
- [26] Y. Nesterov. *Introductory lectures on convex optimization: A basic course*. Springer Science & Business Media, 2006. [2](#)
- [27] R. Johnson and T. Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Proc. Conf. Neural Information Processing Systems*, pages 315–323, 2013. [2](#)
- [28] L. Nguyen, J. Liu, K. Scheinberg, and M. Takáč. SARAH: A novel method for machine learning problems using stochastic recursive gradient. *Proc. Int'l Conf. Machine Learning*, 2018. [2](#)
- [29] L. Nguyen, J. Liu, K. Scheinberg, and M. Takáč. Stochastic recursive gradient algorithm for nonconvex optimization. *arXiv preprint arXiv:1705.07261*, 2017. [2](#)
- [30] C. Fang, C. Li, Z. Lin, and T. Zhang. SPIDER: Near-optimal non-convex optimization via stochastic path integrated differential estimator. *arXiv preprint arXiv:1807.01695*, 2018. [2](#), [4](#), [5](#)
- [31] Z. Wang, K. Ji, Y. Zhou, Y. Liang, and V. Tarokh. Spiderboost: A class of faster variance-reduced algorithms for nonconvex optimization. *arXiv preprint arXiv:1810.10690*, 2018. [2](#)
- [32] J. Zhang, H. Zhang, and S. Sra. R-SPIDER: A fast Riemannian stochastic optimization algorithm with curvature independent rate. *arXiv preprint arXiv:1811.04194*, 2018. [3](#)
- [33] W. Huang, P. Absil, and K. Gallivan. A Riemannian symmetric rank-one trust-region method. *Mathematical Programming*, 150(2):179–216, 2015. [3](#)

- [34] W. Huang, K. Gallivan, and P. Absil. A broyden class of quasi-Newton methods for Riemannian optimization. *SIAM Journal on Optimization*, 25(3):1660–1685, 2015. [3](#)
- [35] B. Polyak. Gradient methods for the minimisation of functionals. *USSR Computational Mathematics and Mathematical Physics*, 3(4):864–878, 1963. [3](#)
- [36] Y. Nesterov and B. Polyak. Cubic regularization of Newton method and its global performance. *Mathematical Programming*, 108(1):177–205, 2006. [3](#), [6](#)
- [37] D. Zhou, P. Xu, and Q. Gu. Stochastic nested variance reduction for nonconvex optimization. *arXiv preprint arXiv:1806.07811*, 2018. [5](#)
- [38] H. Karimi, J. Nutini, and M. 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*, pages 795–811. Springer, 2016. [6](#)
- [39] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM*, 58(3):11, 2011. [8](#)