No-regret Online Learning Algorithm on Riemannian Manifold: Full Information and Bandit Feedback

Abstract

We consider the online convex optimization problem on Riemannian manifolds, where a learner attempts to minimize a sequence of time-varying geodesically convex loss functions on Riemannian manifolds. Though many Euclidean online convex optimization algorithms have been proved useful in a wide range of areas, less attention has been paid on their Riemannian counterparts. In this paper, we introduce our Riemannian Online Gradient Descent (R-OGD) algorithm on Hadamard manifolds for both geodesically convex and strongly geodesically convex loss functions, as well as we introduce our Riemannian Bandit (R-BAN) algorithm on Hadamard homogeneous manifolds for geodesically functions. We establish upper regret bounds of the algorithms with respect to time horizon T, manifold curvature κ, and manifold dimension d, and find a universal lower bound of any algorithm for online convex optimization on Hadamard space. All the regret bounds match the corresponding result in Eucli dean spaces. Finally, numerical experiments of Karcher mean problems validate our theoretical result. connectivity. generated convex loss functions

1 Introduction

The *online convex optimization* plays a central role in finding local minimum of time-varying optimizations. Without assuming prior knowledge of the loss function, the online convex optimization predicts solutions before the loss function is revealed. Consequently, the online convex optimization problem has been widely studied in the past decades in optimization, machine learning, and information system [5, 31, 16, 28, 20].

In this paper, we consider the Riemannian online convex optimization (R-OCO) problem,

$$\min_{x_t \in \mathcal{K} \subset \mathcal{M}} f_t(x_t), t = 1, 2, \ldots, T, \tag{1}$$

where $\mathcal{M}$ is a complete Riemannian manifold equipped with a Riemannian metric g and $\mathcal{K}$ is a geodesically convex (g-convex) subset of the Riemannian manifold $\mathcal{M}$. We assume that $\{f_t\}_{t \in t=1,2,\ldots,T}$ are a sequence of unknown functions and every f_t is a geodesically convex (g-convex) function with sufficient smoothness. The R-OCO problem (1) extends the online convex optimization in Euclidean spaces. Additionally, it also contains many important problems with great interest, e.g., machine learning, online principal component analysis (PCA), dictionary learning [24, 17, 37, 22].

The best way to understand the R-OCO problem (1) is as a learning process of T rounds. At each round $t = 1, 2, 3, \ldots, T$, the online learner chooses a strategy x_t from the g-convex subset $\mathcal{K}$. Later or simultaneously, the adversary (or nature) produces a g-convex loss function $f_t : \mathcal{K} \to \mathbb{R}$ of which we have no prior knowledge. Finally, the learner receives the feedback and suffers the loss $f_t(x_t)$. Generally, there are two types of information feedback in consideration. One is the full information feedback, which means that the entire information of f_t is revealed to the learner; the other is the bandit feedback, which means that only the value $f_t(x_t)$ is revealed. The goal of the R-OCO is to

Preprint. Under review.

minimize the *regret*, defined as

$$R(T) = \sum_{t=1}^{T} f_t(x_t) - \min_{x \in \mathcal{K}} \sum_{t=1}^{T} f_t(x),$$

which measures the difference between the cost by $\{x_t\}_{t=1,\ldots,n}$ and the best-fixed point chosen in hindsight. An algorithm is called no-regret [15], if the regret of the algorithm goes sublinearly function with T.

Some classical methods for solving (1) are to treat the manifold as a subset of an ambient Euclidean space and employ Euclidean constrained optimization techniques. For example, [27] presented an algorithm for the online PCA problem, where the variables were updated in an ambient Euclidean space and then projected onto a manifold. However, most classical approaches may be not sensible in practice. For certain manifolds, the dimension of an embedding Euclidean space can be too high (e.g., the Grassmann manifold [13]) and computing the projection can be much expensive (e.g., the symmetric positive-definite matrices [36]). An alternative approach called *Riemannian optimization* is a mainstream technique to solving optimization on manifolds. The Riemannian optimization makes use of intrinsic geometry of manifolds so that it can solve (1) directly on the manifold as an unconstrained problem, and thus avoiding high dimension embedding and high computing cost for the projection. Moreover, this viewpoint has shown benefits from the g-convexity by which a nonconvex optimization problem can be converted into a g-convex one [7, 11]. Consequently, it is attractive to take a Riemannian approach in our problem (1).

Although there have been many existing algorithms for "offline" manifold optimization problems [2, 35, 6], there is little research about the Riemannian online optimization problem. For example, [34] proposed an online EM algorithm for hidden Markov chains on Hadamard homogeneous spaces, while [26] studied the zeroth-order online optimization problem on Hadamard manifolds by the zeroth-order oracle in [25], and then derived upper bounds on the tracking error and regret asymptotically.

Contribution Motivated by the above, this paper aims to design no-regret algorithms for the R-OCO problem in both full information feedback and bandit feedback. The contribution of this paper is summarized as follows:

- We propose a Riemannian online gradient algorithm (R-OGD) for the R-OCO problem in the full information feedback, and then analyze the upper regret bounds of the R-OGD algorithm for g-convex and strongly g-convex functions, which coincides with the regret bounds in Euclidean spaces [38].
- We introduce a Riemannian bandit algorithm (R-BAN) for the R-OCO problem in the bandit feedback and then establish a upper regret bound for g-convex functions are the same as the Euclidean regret in [18]. During the analysis, we develop a key technique to analyze the "smoothed" function over homogeneous Hadamard manifolds, which can be applied to Riemannian optimization on manifolds with symmetry and beyond.
- We achieve a universal lower regret bound of R-OCO algorithms by defining the notion of Riemannian online Busemann optimization on arbitrary Hadamard manifolds, which firstly extend the result in [1].

Related Work The first work about the Euclidean online convex optimization was discussed in [38]. Inspired by the gradient descent method, [38] proposed the online gradient descent algorithm (OGD) of which the upper regret bound of was proved to be $\mathcal{O}\big(\sqrt{T}\big)$. [21] proceeded with the study of the OGD algorithm and established an upper regret bound $\mathcal{O}\big(\log T\big)$ for strongly convex functions. In addition, [1] gave a universal lower bound of $\mathcal{O}\big(\sqrt{T}\big)$ for online convex algorithms, which means that the bounds in [38] and [21] are essentially optimal. In the bandit setting, [18] provided a detailed exposition of a one-point bandit algorithm. By modifying the gradient in the OGD algorithm to a randomized estimator, the upper regret bound attained $\mathcal{O}\big(T^{\frac{3}{4}}\big)$ and $\mathcal{O}\big(T^{\frac{2}{3}}\big)$ for convex loss functions and strongly convex loss functions respectively. Extending the one-point bandit algorithm, [4] developed a multi-point bandit algorithm and presented upper regret bounds $\mathcal{O}\big(\sqrt{T}\big)$ and $\mathcal{O}\big(\log T\big)$ for convex and strongly convex loss functions. In this work, we design Riemannian

online algorithms in full information feedback and one-point bandit feedback setting, which recover the regret bounds ([38, 21, 18, 1]) in Euclidean spaces.

Riemannian optimization has drawn much research attention in the past decades. Many basic algorithms in Euclidean spaces such as gradient descent method, Newton's method, and trust region methods have been adapted into a Riemannian setting [3, 6, 30, 32, 29, 23]. Some work of Riemannian stochastic optimization are intended to deal with time-varying optimization problems[12, 35, 36, 34]. Among them, [35] has provided the first global complexity analysis for the R-SGD on geodesically convex problems over Hadamard manifolds and [34] has proposed an online EM algorithm on Hadamard homogeneous spaces. However, all algorithms mentioned above assume prior knowledge of loss functions, which is different from the framework of our R-OCO problem. The work about the R-OCO problem is fairly limited. To the best of our knowledge, only very recently [26] has proposed a zeroth-order online algorithm on Hadamard manifolds, but their analysis on regret was merely given asymptotically. In this paper, we propose algorithms that rely on no prior knowledge of loss functions and analyze them with explicit regret bounds with the time horizon T.

2 Preliminaries

Riemannian manifolds A *Riemannian manifold* is a smooth manifold $\mathcal{M}$ equipped with a metric tensor g (or called Riemannian metric), which defines an inner product $\langle \cdot, \cdot \rangle_x$ in every tangent space $T_x\mathcal{M}$. The Riemannian metric g brings a distance structure on $\mathcal{M}$. A curve is a *geodesic* if it locally minimizes the distance, which is an analog of the straight line in Euclidean spaces. On Riemannian manifolds, a geodesic is uniquely determined by its start point and initial tangent vector. In this way, the *exponential map* $\exp_x : T_x\mathcal{M} \to \mathcal{M}$ is defined by mapping a vector $X \in T_p\mathcal{M}$ to $\gamma(1) \in \mathcal{M}$ for the geodesic γ such that $\gamma(0) = x$ and $\dot{\gamma}(0) = X$.

The curvature tensor reflects the geometry of manifolds. We focus on one form of curvature called *sectional curvature* which is the Gauss curvature of a two-dimensional submanifold. Following from the work in [35], we consider the *Hadamard manifold*, which is a simply connected and complete manifold with nonpositive curvature. The Cartan-Hadamard theorem [10] shows that the exponential map has an inverse $\exp_x^{-1}()$ globally on Hadamard manifolds and thus the distance $d(x, y)$ can be expressed as $\| \exp_x^{-1}(y)\|$.

Isometries of Riemannian manifolds has been widely studied in differential geometry[10, 9]. An isometry $\phi : \mathcal{M} \to \mathcal{M}$ is a diffeomorphism preserving distance, i.e., $d(x, y) = d(\phi(x), \phi(y))$ for all $x, y \in \mathcal{M}$. It is remarked that all isometries of a Riemannian manifold form a Lie group G. A Riemannian manifold is called *homogeneous manifold* if the group of isometries G acts on $\mathcal{M}$ transitively, i.e., for any points $x, y \in \mathcal{M}$ there exists an isometry such that $\phi(x) = y$.

Function Classes We now define some special classes of functions that are useful in Riemannian optimization. A set $\mathcal{K}$ is called *geodesically convex* (g-convex) if, for any points $x, y \in \mathcal{K}$, there admits a geodesic γ contained in $\mathcal{K}$ such that $\gamma(0) = x$ and $\gamma(1) = y$. A function $f : \mathcal{K} \to \mathbb{R}$ is called *geodesically convex* (or g-convex) if for any geodesic $\gamma : [0, 1] \to \mathcal{M}$, it satisfies

$$f(\gamma(t)) \leq (1 - t)f(\gamma(0)) + tf(\gamma(1)).$$

If f is differentiable, the definition is equivalent to the following condition

$$f(y) \geq f(x) + \langle \nabla f(x), \exp_x^{-1}(y)\rangle, \forall x, y \in \mathcal{M}.$$

A function $f : \mathcal{M} \to \mathbb{R}$ is *geodesically μ-strongly convex* (or μ-strongly g-convex) if there exists a constant $\mu > 0$, for any $x, y \in \mathcal{M}$, it satisfies that

$$f(y) \geq f(x) + \langle \nabla f(x), \exp_x^{-1}(y)\rangle + \frac{\mu}{2} d^2(x, y).$$

We called a function is *geodesically L-Lipschitz* (g-L-Lipschitz) if there exists a constant $L > 0$, for any $x, y \in \mathcal{M}$, it satisfies

$$|f(y) - f(x)| \leq L \cdot d(x, y).$$

If f is differentiable, the definition is equivalent to the condition that

$$\|\nabla f(x)\| \leq L, \forall x \in \mathcal{M}.$$

3 Riemannian Online Convex Optimization with Full Information Feedback

This section is devoted to the study of the R-OCO problem in the full information feedback. In this section, we first propose our R-OGD algorithm, and then analyze the upper regret bounds of R-OGD for both g-convex functions and strongly g-convex cases. In addition, a universal lower regret bound is presented to illustrate that the regret bound of our R-OGD algorithm is tight up to a constant.

3.1 Riemannian Online Gradient Algorithm

In the full information setting, we consider the following assumptions, which were used in the existing literature of Euclidean online convex optimization and differential geometry, (e.g., [38, 6, 23]).

Assmuption 1. *There exists a $x^\star \in \mathcal{M}$ such that $x^\star = \arg\min \sum_{t=1}^{T} f_t(x)$.*

Assmuption 2. *$(\mathcal{M}, g)$ is a Hadamard manifold with sectional curvature lower bounded by a constant $-\kappa$ ($\kappa > 0$).*

Remark 1. The Hadamard manifold plays an important role in Riemannian geometry (e.g. [19]). Some basic spaces, such as the Euclidean space $\mathbb{R}^n$, the Hyperbolic space H^n, the positive symmetry definite matrix manifold $\mathbb{S}_{++}$, are all Hadamard manifolds. Consequently, the assumption of the Hadamard manifold is widely used in Riemannian optimization (e.g.[6, 23]).

Assmuption 3. *The g-convex set $\mathcal{K}$ is a bounded set with the diameter D, i.e.,*

$$d(x, y) \leq D, \forall x, y \in \mathcal{K}.$$

Assmuption 4. *For all $t \in [T]$, f_t are differentiable and g-L-Lipschitz.*

We now propose our Riemannian online gradient descent algorithm (R-OGD) in Algorithm 1, where we use the exponential map to replace the vector addition in [38].

Algorithm 1: Riemannian Online Gradient Descent Algorithm

Input: Manifold $\mathcal{M}$, time T, step size $\{\eta_t\}$
Output: $\{x_i\}_{t\in[T]}$
for $t = 1$ *to* T *do* **do**
  Play x_t and observe the function f_t;
  Update

$$x_{t+1} = \mathcal{P}_{\mathcal{K}}(\exp_{x_t}(-\nabla f_t(x_t))),$$

  where $\exp_{x_t}(\cdot)$ is the exponential map at x_t and $\mathcal{P}_{\mathcal{K}}$ is the Riemannian projection mapping of x onto $\mathcal{K}$, i.e., $\mathcal{P}_{\mathcal{K}}(x) := \arg\min_{y\in\mathcal{K}} d(x, y)$;
  Return x_{t+1}, and suffer from the loss $f_t(x_t)$;
end

3.2 Upper Regret Bound Analysis

Next, we analyze regret bounds of the R-OGD algorithm. Theorems 1 and 2 give analyses of the upper regret bounds for g-convex loss functions and strongly g-convex functions respectively. We use the symbol $\zeta(\kappa, d)$ to denote the function $\frac{\sqrt{\kappa}\cdot d}{\tanh\left(\sqrt{\kappa}\cdot d\right)}$. By direct observation, ζ is a increasing function of the variables κ and d when $\kappa d \geq 0$.

Theorem 1 (Convex Cases)**.** *Suppose that assumptions 1-4 hold. If every f_t is g-convex, then the R-OGD algorithm with step sizes $\{\eta_t = \frac{D}{L\sqrt{\zeta(\kappa,D)t}}\}$ guarantees the following for all $T \geq 1$:*

$$R(T) = \sum_{t=1}^{T} f_t(x_t) - \sum_{t=1}^{T} f(x^\star) \leq \frac{3}{2} DL\sqrt{\zeta(\kappa, D)\cdot T}$$

Theorem 2 (Strongly-convex Cases)**.** *Suppose that assumptions 1-4 hold. If every f_t is μ-strongly g-convex, then the R-OGD algorithm with step sizes $\{\eta_t = \frac{1}{\mu t}\}$ guarantees the following for all $T \geq 1$:*

$$R(T) = \sum_{t=1}^{T} f_t(x_t) - \sum_{t=1}^{T} f(x^\star) \leq \frac{L^2\zeta(\kappa, D)}{2\mu}(1 + \log T).$$

The proofs of Theorems 1 and 2 are in the appendix. The major challenge in proving Theorems 1 and 2 is that there is no vector space structure on Riemannian manifolds. Thank to the trigonometric distance bound proposed in [35], we manage to obtain the regret $\mathcal{O}\big(\sqrt{T}\big)$ and $\mathcal{O}\big(\log T\big)$ for g-convex loss functions and strongly g-convex loss functions. By slightly moving κ to zero, the results recover the standard results in [38] and [21].

This result also shows the influence of curvature on the regret bounds. Since $\zeta(\kappa, d)$ is a increasing function of κ, the upper regret bounds in the R-OGD algorithm is higher than those in Euclidean spaces and the increase of κ raises the upper regret bound. Therefore, proper Riemannian metric should be chosen in the optimization to avoid high sectional curvature bound.

3.3 Lower Regret Bound Analysis

This section is intended to answer the question of whether there exists an algorithm that attains a tighter regret bound than $\mathcal{O}\big(\sqrt{T}\big)$ for g-convex functions. In this section, we answer the question negatively by constructing the Riemannian online Busemann optimization.

In Euclidean space, the lower bound is attained in [1] by analyzing an online convex optimization with linear loss. However, this method is not well adapted to Riemannian manifolds as there is no linear functions on Riemannian manifolds. Under this circumstance, we introduce the *Riemannian online Busemann optimization*. Let $\mathcal{M}$ be a Hadamard manifold, p be a point on the manifold and $\gamma : \mathbb{R} \to \mathcal{M}$ be a geodesic with $\|\dot{\gamma}(0) = 1\|$ $\gamma(0) = p$. Then we define the online Riemannian online Busemann optimization on $\mathcal{M}$ as a learning process lasting T rounds. In each round, the learner's choice comes from the set $\mathcal{K} = B_1(p)$, and the loss function f_t is randomly and uniformly chosen in set $\{f_+, f_-\}$, where f_+ and f_- is the Busemann functions [8] with γ_+ and γ_- where $\gamma_\pm(t) = \gamma(\pm t), t \in (0, \infty]$ are two geodesic rays related to γ.

Riemannian online Busemann optimization

Data: $\mathcal{K}, \{f_+, f_-\}$
for t= 1 *to* T *do* **do**
  The learner chooses (predicts) $x_t \in \mathcal{K}$.
  The loss function f_t is randomly and uniformly chosen in set $\{f_+, f_-\}$
end
The learner suffers regret

$$R(T) = \sum_{i=1}^{T} f_t(x_t) - \min_{x \in \mathcal{K}} f_t(x).$$

Theorem 3. *Suppose that assumptions 1-4 hold. Then for any Hadamard manifold $\mathcal{M}$, the expected regret of the Riemannian online Busemann optimization $E[R(T)]$ satisfies:*

$$E[R(T)] \geq \frac{1}{\sqrt{2}}\sqrt{T},$$

which implies the Riemannian online convex optimization incurs the regret $\mathcal{O}\big(DG\sqrt{T}\big)$ for any possible strategy in the worst case.

The proof of Theorem 3 is in the appendix. The result illustrates that, as the same as that in the Euclidean space, the regret of any Riemannian online algorithm can not reach less than $\mathcal{O}\big(\sqrt{T}\big)$ in the worst case. Moreover, the result concludes that the regret of the R-OGD algorithm is tight up to a constant.

4 Riemannian Online Convex Optimization with Bandit Feedback

In this section, we consider the Riemannian online convex optimization with the one-point bandit feedback. We first present the Riemannian bandit algorithm (R-BAN) on Hadamard homogeneous manifolds and then give analysis on the (expected) upper regret bound for our algorithm. In the following of this section, for any point x, we denote $B_\delta(x)$ as the ball centered at x with radius δ, and $V_\delta(x)$ as its n-dim volume. Correspondingly, we denote $S_\delta(x)$ as the sphere centered at x with radius δ, and $S_\delta(x)$ as its $(n-1)$-dim volume.

4.1 Riemannian Bandit Algorithm

In the bandit setting, assumptions 2-4 should be slightly modified to adapt the bandit case. All of them were widely used in bandit optimization and Riemannian geometry, see, e.g., [18, 4].

Assmuption 5. *$\mathcal{M}$ is an n-dimensional homogeneous Hadamard manifold with sectional curvature lower bounded by a constant $-\kappa$ ($\kappa > 0$).*

Remark 2. The homogeneous Hadamard manifold is widely researched in differential geometry (e.g.,[9, 10]). The property of homogeneity has received much attention in machine learning [33, 34, 14]. It is seen that many manifolds that are often considered in Riemannian optimization, such as the Euclidean space $\mathbb{R}^n$, the Hyperbolic space H^n, the positive symmetry definite matrix manifold $\mathbb{S}_{++}$, are Hadamard homogeneous manifolds. Therefore, we consider the assumption of the homogeneous Hadamard manifold to be reasonably valid.

Note that on Hadamard homogeneous manifolds, the volume and surface area of a ball is only related to the radius but not to the center of the ball, thus we denote V_δ as the volume of $B_\delta(x)$ and S_δ as $(n-1)$-dim volume (surface area) of $S_\delta(x)$ for all x in the manifold $\mathcal{M}$.

Assmuption 6. *Suppose there exists a interior point $p \in \mathcal{K}$ such that the set $\mathcal{K}$ contains a ball with radius r centered at p, and is also contained in a ball with radius D, i.e.,*

$$B_r(p) \subset \mathcal{K} \subset B_D(p).$$

Assmuption 7. *For all $t \in [T]$, f_t are differentiable, g-L-Lipschitz and bounded by C.*

Inspired by the Euclidean bandit algorithm, we replace the gradient $\nabla f_t(x_t)$ with a randomized estimator g_t and propose our Riemannian bandit optimization algorithm (R-BAN) in Algorithm 2 over Hadamard homogeneous manifolds.

Algorithm 2: Riemannian Bandit Algorithm

Input: Manifold $\mathcal{M}$, time T, step size ν, parameter δ, α.
Output: Sequence $\{x_i\}_{t\in[T]}$
for $t = 1$ *to* T *do* **do**
  Pick x_t uniformly from $S_\delta(y_t)$;
  Play x_t and observe $f_t(x_t)$;
  Construct the gradient estimator

$$g_t = f_t(x_t)\frac{\exp_{y_t}^{-1}(x_t)}{\|\exp_{y_t}^{-1}(x_t)\|};$$

  Update with the rule

$$y_{t+1} = P_{(1-\alpha)K}(\exp_{y_t}(-\nu g_t)),$$

  where we denote $P_{(1-\alpha)K}$ as the projection mapping onto the shrinking set $\mathcal{P}(1-\alpha)\mathcal{K} = \{\exp_p((1-\alpha)u) | u = \exp_p^{-1}(x), x \in \mathcal{K}\}$.
  Return x_t and suffer from the loss $f_t(x_t)$;
end

4.2 Difficulties

Since the proposed algorithm 2 is a straightforward extension of the Euclidean bandit algorithm in [18], it is worth reviewing the analysis in [18]. In Euclidean setting, we uniformly choose x_t on the $S_\delta x_t$ and update by the rule

$$\begin{cases} \tilde{g}_t = \frac{n}{\delta} f(y_t) \frac{y_t - x_t}{\|y_t - x_t\|}, \\ x_{t+1} = \mathcal{P}_{(1-\alpha)\mathcal{K}}(x_t - \nu g_t)), \end{cases} \tag{2}$$

The basis idea for the analysis is to introduce the *smoothed loss function*

$$\hat{f}_t^E(x_t) = E_{u\in B_\delta(x)}[f(u)] = \frac{1}{V_\delta}\int_{B_\delta(x_t)} f_t(u)du,$$

where $\hat{f}_t^E$ is a convex approximation of f_t when δ is small. It is shown that $\frac{S_\delta}{V_\delta}\tilde{g}_t$ is an unbiased estimator of the gradient $\nabla \hat{f}_t^E(x_t)$, hence the bandit algorithm is actually a expected gradient descent method with the loss function $\hat{f}_t^E$. In this way, a regret bound of the bandit algorithm is established by [18].

Back to the Riemannian case, we attempt to generalize the analysis of [18] in parallel by defining the "Riemannian version" of the smoothed loss function,

$$\hat{f}_t(x_t) = E_{u\in B_\delta(x)}[f(u)] = \frac{1}{V_\delta}\int_{B_\delta(x_t)} f_t(u)\omega,$$

where ω is the volume element with the respect to the metric g. However, carrying out such definition is not enough for directly generalizing further analysis on regret bounds, due to at least two difficulties.

The gradient is hard to compute. Computing the gradient of $\hat{f}_t$ is quite different from that in Euclidean spaces, as the absence of the commutativity of the derivative operator ∇ and the integration operator $\int_{B_\delta(x)}$. In Euclidean spaces, the derivative operator ∇ commutes with the integration operator $\int_{B_\delta(x)}$. Accordingly, for the Euclidean smoothed loss function $\hat{f}_t^E$,

$$\nabla \hat{f}_t^E(x_t) = \frac{1}{V_\delta}\nabla \int_{B_\delta(x_t)} f_t(u)du = \frac{1}{V_\delta}\int_{B_\delta(x_t)} \nabla f_t(u)du. \tag{3}$$

which derives the fact $\frac{S_\delta}{V_\delta}E[\tilde{g}_t] = \nabla \hat{f}_t^E(x_t)$. However, on Riemannian manifolds the derivative operator ∇ does not commute with the integration operator $\int_{B_\delta(x)}$. Consequently, the equation (3) fails to hold for Riemannian manifolds, making it is hard to compute the gradient.

The convexity may not hold. Another essential challenge for regret analysis is the convexity of $\hat{f}_t$. In Euclidean spaces, one can easily conclude the convexity of $\hat{f}_t$, because the Hessian of $\hat{f}_t(x_t)$ is a local average of the Hessian of f_t. However, the convexity may not hold for Riemannian manifolds.

4.3 Solution to the Difficulties and Regret Analysis

To address the first difficulty, we propose a key technique to analyze the derivative of local integration by introducing the Killing vector field. With the help of this technique, we manage to compute the gradient of $\hat{f}_t$ in Lemma 1.

Lemma 1. *Let $\mathcal{M}$ be a homogeneous Hadamard manifold. Suppose that f is a C^1 function on $\mathcal{M}$, Then the smoothed function $\hat{f}(x) = \frac{1}{V_\delta}\int_{B_\delta(x)} f\omega$ satisfies,*

1. $\nabla \hat{f}(x) = \frac{1}{V_\delta}\int_{S_\delta x} f(u)\frac{\exp_x^{-1}(u)}{\|\exp_x^{-1}(u)\|} = \frac{S_\delta}{V_\delta}E_{u\in S_\delta(x)}\left[f(u)\frac{\exp_x^{-1}(u)}{\|\exp_x^{-1}(u)\|}\right], \forall x \in \mathcal{M}$;
2. *if $|f(x)| < C$, then there exist a number $\epsilon(n)$ such that*

$$\|\nabla \hat{f}(x)\| \le \frac{S_\delta}{V_\delta}C \le \frac{n}{\delta}C + \frac{(n^2-n)}{3n+6}\kappa\delta C$$

holds for all $0 < \delta < \epsilon(n)$.

Remark 3. The proof of Lemma 1 can be found in the appendix. The first part of the lemma proceeds the result in [18], while the second part gives us an easy-to-compute bound of the gradient. From the proof, we develop a key technique that transform derivation of integration on $B_\delta(x)$ to integration of derivative of corresponding Killing vector field on $B_\delta(x)$, i.e.,

$$\nabla_X\Big(\int_{B_\delta(x)} f(u)\omega\Big) = \int_{B_\delta(x)} \nabla_\eta f(u)\omega, \tag{4}$$

where η is a Killing field such that $\eta(x) = X$. This technique does not rely on the curvature and other specific manifold structures, hence it can be a basic tool for optimization on homogeneous spaces and maybe in broader areas.

To overcome the second difficulty, we notice that though the function $\hat{f}_t$ may not be g-convex, but it is very closely to be g-convex. We conclude it in Lemma 2.

Lemma 2. *Suppose that $\mathcal{M}$ is a Hadamard homogeneous manifold and $\mathcal{K}$ is a convex and bounded set of $\mathcal{M}$. Then there exists a number $\rho \geq 0$ only related to the set $\mathcal{K}$ such that for any g-convex and g-L-Lipschitz function f, the smoothed function $\hat{f}$ satisfies*

$$\hat{f}(y) - \hat{f}(x) - \langle \nabla \hat{f}(x), \exp_x^{-1}(y) \rangle \geq -2\rho\delta L, \forall x, y \in \mathcal{K}$$

The proof of Lemma 2 is in the appendix. It is worth discussing the constant ρ which describes how close a smoothed function is from being g-convex. Notice that ρ does not related to the function $\hat{f}$, it is shown that the gap from g-convexity is uniform for all smoothed functions. Moreover, if the manifold $\mathcal{M}$ reduces to a Euclidean space, then we can find $\rho = 0$ and $\hat{f}$ is convex, which coincides the result in Euclidean spaces. Additionally, though the value of ρ is generally difficult to calculate, our algorithm analysis and parameter selection do not depend on the specific value of ρ (see Theorem 4).

After surmounting the difficulties above, we now carry out the analysis of the expected regret bounds of Algorithm 2.

Theorem 4. *Suppose that assumptions 1 and 5-7 holds. We assume that every f_t is g-convex. Denote $B = \frac{n^2-n}{3n+6}\kappa$. For any $T > \left(\frac{CDnr\sqrt{\zeta(\kappa,D)}}{BCDr\sqrt{\zeta(\kappa,D)}+3Lr+2C}\right)^2/\epsilon(n)^4$, If we take $\alpha = \frac{\delta}{r}$, $\nu = \frac{D}{C\sqrt{\zeta(\kappa,D)T}}$, $\delta = T^{-\frac{1}{4}}\sqrt{\frac{CDnr\sqrt{\zeta(\kappa,D)}}{BCDr\sqrt{\zeta(\kappa,D)}+3Lr+2C}}$, then the expected regret of algorithm 2 is upper bounded by*

$$E\Big[\sum_{t=1}^{T} f_t(x_t)\Big] - \min_{x\in\mathcal{K}} \sum_{t=1}^{T} f_t(x) \leq$$
$$2T^{\frac{3}{4}}\sqrt{nCD\sqrt{\zeta(\kappa,D)}}\cdot\Big(\sqrt{(BCD\sqrt{\zeta(\kappa,D)}+3L+2C/r)} + 2\frac{\rho L}{\sqrt{(BCD\sqrt{\zeta(\kappa,D)}}}\Big).$$

The proof of Theroem 4 can be seen in the appendix. Theorem 4 shows that the regret of the Riemannian bandit algorithm achieves $\mathcal{O}\big(T^{\frac{3}{4}}\big)$ for g-convex loss functions on homogeneous Hadamard manifolds, which is same as the regret bounds in Euclidean spaces [18].

Compare with the result in [26], our work is different in two aspects. First, in [26] the tracking error and regrets bounds are given asymptotically, while in Theorem 4, the regret bound is expressed explicitly with the time horizon T. Second, to estimate the gradient, [26] use a biased estimator of the gradient, while in our bandit algorithm the gradient estimator g_t of the smoothed function is unbiased.

5 Numerical Experiment

References

[1] Jacob Abernethy, Peter L Bartlett, Alexander Rakhlin, and Ambuj Tewari. Optimal strategies and minimax lower bounds for online convex games. 2008.

[2] P-A Absil, Christopher G Baker, and Kyle A Gallivan. Trust-region methods on riemannian manifolds. *Foundations of Computational Mathematics*, 7(3):303–330, 2007.

[3] P-A Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization algorithms on matrix manifolds*. Princeton University Press, 2009.

[4] Alekh Agarwal, Ofer Dekel, and Lin Xiao. Optimal algorithms for online convex optimization with multi-point bandit feedback. In *Proceedings of the 23rd Annual Conference on Learning Theory (COLT)*, June 2010.

[5] Shmuel Agmon. The relaxation method for linear inequalities. *Canadian Journal of Mathematics*, 6:382–392, 1954.

[6] Kwangjun Ahn and Suvrit Sra. From nesterov's estimate sequence to riemannian acceleration. *arXiv preprint arXiv:2001.08876*, 2020.

[7] Zeyuan Allen-Zhu, Ankit Garg, Yuanzhi Li, Rafael Oliveira, and Avi Wigderson. Operator scaling via geodesically convex optimization, invariant theory and polynomial identity testing. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 172–181, 2018.

[8] Werner Ballmann. *Lectures on spaces of nonpositive curvature*, volume 25. Birkhäuser, 2012.

[9] Valerii Berestovskii and Yurii Nikonorov. *Riemannian manifolds and homogeneous geodesics*. Springer, 2020.

[10] Marcel Berger. *Geometry i*. Springer Science & Business Media, 2009.

[11] Dario A Bini and Bruno Iannazzo. Computing the karcher mean of symmetric positive definite matrices. *Linear Algebra and its Applications*, 438(4):1700–1710, 2013.

[12] Silvere Bonnabel. Stochastic gradient descent on riemannian manifolds. *IEEE Transactions on Automatic Control*, 58(9):2217–2229, 2013.

[13] Nicolas Boumal and P-A Absil. Low-rank matrix completion via preconditioned optimization on the grassmann manifold. *Linear Algebra and its Applications*, 475:200–239, 2015.

[14] Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges, 2021.

[15] Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.

[16] Ofer Dekel, Shai Shalev-Shwartz, and Yoram Singer. The power of selective memory: Self-bounded learning of prediction suffix trees. In *Proceedings of the 17th International Conference on Neural Information Processing Systems*, pages 345–352, 2004.

[17] Jiashi Feng, Huan Xu, and Shuicheng Yan. Online robust pca via stochastic optimization. In *Advances in neural information processing systems*, pages 404–412. Citeseer, 2013.

[18] Abraham D. Flaxman, Adam Tauman Kalai, and H. Brendan McMahan. Online convex optimization in the bandit setting: Gradient descent without a gradient. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, page 385–394, USA, 2005. Society for Industrial and Applied Mathematics.

[19] Mohammad Ghomi and Joel Spruck. Total curvature and the isoperimetric inequality in cartan-hadamard manifolds, 2021.

[20] Elad Hazan. Introduction to online convex optimization. *arXiv preprint arXiv:1909.05207*, 2019.

[21] Elad Hazan, Adam Kalai, Satyen Kale, and Amit Agarwal. Logarithmic regret algorithms for online convex optimization. In *International Conference on Computational Learning Theory*, pages 499–513. Springer, 2006.

[22] Jiang Hu, Xin Liu, Zai-Wen Wen, and Ya-Xiang Yuan. A brief introduction to manifold optimization. *Journal of the Operations Research Society of China*, 8(2):199–248, 2020.

[23] Wen Huang, P-A Absil, and Kyle A Gallivan. A riemannian symmetric rank-one trust-region method. *Mathematical Programming*, 150(2):179–216, 2015.

[24] Kuang-Chih Lee and David Kriegman. Online learning of probabilistic appearance manifolds for video-based recognition and tracking. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 852–859. IEEE, 2005.

[25] Jiaxiang Li, Krishnakumar Balasubramanian, and Shiqian Ma. Stochastic zeroth-order riemannian derivative estimation and optimization, 2021.

[26] Alejandro I Maass, Chris Manzie, Dragan Nesic, Jonathan H Manton, and Iman Shames. Online zeroth-order optimisation on riemannian manifolds. *arXiv preprint arXiv:2010.00211*, 2020.

[27] Jiazhong Nie, Wojciech Kotłowski, and Manfred K Warmuth. Online pca with optimal regret. *The Journal of Machine Learning Research*, 17(1):6022–6070, 2016.

[28] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '14, page 701–710, New York, NY, USA, 2014. Association for Computing Machinery.

[29] Chunhong Qi, Kyle A Gallivan, and P-A Absil. Riemannian bfgs algorithm with applications. In *Recent advances in optimization and its applications in engineering*, pages 183–192. Springer, 2010.

[30] Wolfgang Ring and Benedikt Wirth. Optimization methods on riemannian manifolds and their application to shape space. *SIAM Journal on Optimization*, 22(2):596–627, 2012.

[31] Frank Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386, 1958.

[32] Berkant Savas and Lek-Heng Lim. Quasi-newton methods on grassmannians and multilinear approximations of tensors. *SIAM Journal on Scientific Computing*, 32(6):3352–3393, 2010.

[33] Hao Tang, Zhiao Huang, Jiayuan Gu, Bao-Liang Lu, and Hao Su. Towards scale-invariant graph-related problem solving by iterative homogeneous gnns. *Advances in Neural Information Processing Systems*, 33, 2020.

[34] Quinten Tupker, Salem Said, and Cyrus Mostajeran. Online learning of riemannian hidden markov models in homogeneous hadamard spaces. *arXiv preprint arXiv:2102.07771*, 2021.

[35] Hongyi Zhang and Suvrit Sra. First-order methods for geodesically convex optimization. In *Conference on Learning Theory*, pages 1617–1638, 2016.

[36] Jingzhao Zhang, Hongyi Zhang, and Suvrit Sra. R-spider: A fast riemannian stochastic optimization algorithm with curvature independent rate. *arXiv preprint arXiv:1811.04194*, 2018.

[37] Shengping Zhang, Shiva Kasiviswanathan, Pong Yuen, and Mehrtash Harandi. Online dictionary learning on symmetric positive definite manifolds with vision applications. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.

[38] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning*, ICML'03, page 928–935. AAAI Press, 2003.