

Riemannian Adaptive Optimization Algorithm and Its Application to Natural Language Processing

Hiroyuki Sakai, Hideaki Iiduka

Abstract—This paper proposes a Riemannian adaptive optimization algorithm to optimize the parameters of deep neural networks. The algorithm is an extension of both AMSGrad in Euclidean space and RAMSGrad on a Riemannian manifold. The algorithm helps to resolve two issues affecting RAMSGrad. The first is that it can solve the Riemannian stochastic optimization problem directly, in contrast to RAMSGrad which only achieves a low regret. The other is that it can use constant learning rates, which makes it implementable in practice. Additionally, we apply the proposed algorithm to Poincaré embeddings that embed the transitive closure of the WordNet nouns into the Poincaré ball model of hyperbolic space. Numerical experiments show that regardless of the initial value of the learning rate, our algorithm stably converges to the optimal solution and converges faster than the existing algorithms.

Index Terms—Riemannian optimization, RSGD, RAdaGrad, RAdam, Riemannian adaptive optimization algorithm, RAMSGrad, Natural language processing, Poincaré embeddings, Hyperbolic space

I. INTRODUCTION

R IEMANNIAN optimization has attracted a great deal of attention [1]–[3] in light of developments in machine learning and deep learning. This paper focuses on Riemannian adaptive optimization algorithms for solving an optimization problem on a Riemannian manifold. In the field of machine learning, there is an important example of the Riemannian optimization problem. Nickel and Kiela [4] proposed Poincaré embeddings, which embed hierarchical representations of symbolic data (e.g., text, graph data) into the Poincaré ball model of hyperbolic space. In fact, experiments on transitive closure of the WordNet noun hierarchy showed that embeddings into a 5-dimensional Poincaré ball are better than embeddings into a 200-dimensional Euclidean space. Since the Poincaré ball has a Riemannian manifold structure, the problem of finding Poincaré embeddings should be considered to be a Riemannian optimization problem.

Bonnabel [5] proposed Riemannian stochastic gradient descent (RSGD), the most basic Riemannian stochastic optimization algorithm. RSGD is a simple algorithm, but its slow convergence is problematic. In [6], Sato, Kasai, and Mishra proposed the Riemannian stochastic variance reduced gradient (RSVRG) algorithm and gave a convergence analysis under some natural assumptions. RSVRG converges to an optimal

solution faster than RSGD; however, RSVRG needs to calculate the full gradient every few iterations. In Euclidean space, adaptive optimization algorithms, such as AdaGrad [7], Adam [8, Algorithm 1], Adadelta [9], and AMSGrad [10, Algorithm 2], [11, Algorithm 1], are widely used for training deep neural networks. However, these adaptive algorithms cannot be naturally extended to general Riemannian manifolds, due to the absence of a canonical coordinate system. Therefore, special measures are required to extend the adaptive algorithms to Riemannian manifolds. For instance, Kasai, Jawanpuria, and Mishra [12] proposed adaptive stochastic gradient algorithms on Riemannian matrix manifolds by adapting the row, and column subspaces of gradients.

In the particular case of a product of Riemannian manifolds, Bécigneul and Ganea [13] proposed Riemannian AMSGrad (RAMSGrad) by regarding each component of the product Riemannian manifold as a coordinate component in Euclidean space. However, their convergence analysis had two points requiring improvement. First, they only performed a regret minimization (Theorem III.2) and did not solve the Riemannian optimization problem. Second, they did a convergence analysis with only a diminishing learning rate; i.e., they did not perform a convergence analysis with a constant learning rate. Since diminishing learning rates are approximately zero after a large number of iterations, algorithms that use them are not implementable in practice. In contrast, a constant learning rate does not cause this problem.

The motivation of this work is to identify whether or not RAMSGrad can be modified in such a way that it can be applied to Riemannian optimization from the viewpoints of both theory and practice. The theoretical motivation is to show that the modified RAMSGrad can solve directly the Riemannian optimization problem, while the practical motivation is to show that it can be applied to important problems in natural language processing and principal component analysis.

Motivated by the above discussion, we propose modified RAMSGrad (Algorithm 1), which is an extension of RAMSGrad, to solve the Riemannian optimization problem (Problem II.1). In addition, we give a convergence analysis (Theorem III.1) valid for both a constant learning rate (Corollary III.1) and diminishing learning rate (Corollary III.2). The analysis leads to the finding that the proposed algorithm can solve directly the Riemannian optimization problem. In particular, we emphasize that the proposed algorithm can use a constant learning rate to solve the problem (Corollary III.1), in contrast to the previous result [13] in which RAMSGrad with a diminishing learning rate only performed a regret minimization (see Subsection III-B for comparisons of the proposed algorithm

H. Sakai is with the Department of Computer Science, Meiji University, Kanagawa 214-8571, Japan (e-mail: sakai0815@cs.meiji.ac.jp).

H. Iiduka is with the Department of Computer Science, Meiji University, Kanagawa 214-8571, Japan (e-mail: iiduka@cs.meiji.ac.jp).

This work was supported by JSPS KAKENHI Grant Number JP18K11184.

with RAMSGrad). In numerical experiments, we apply the proposed algorithm to Poincaré embeddings and compare it with RSGD, Riemannian AdaGrad (RAdaGrad) [13, Section 3.2], and Riemannian Adam (RAdam) [13, Section 4] (Section IV). We show that it converges to the optimal solution faster than the existing algorithms and that it minimizes the objective function regardless of the initial learning rate. In particular, we show that the proposed algorithm with a constant learning rate is a good way of embedding the WordNet mammals subtree into a Poincaré ball. We also applied the algorithm to principal component analysis and found that the choice between using a constant or a diminishing learning rate depends on the dataset. These numerical comparisons lead to the finding that the proposed algorithm is good for solving important problems in natural language processing and principal component analysis.

This paper is organized as follows. Section II gives the mathematical preliminaries and states the main problem. Section III describes the modified RAMSGrad and gives its convergence analysis. Section IV numerically compares the performance of the proposed learning algorithms with the existing algorithms. Section V concludes the paper with a brief summary.

II. MATHEMATICAL PRELIMINARIES

A. Definitions, assumptions, and main problem

Let M be a Riemannian manifold. An exponential map at $x \in M$, written as $\exp_x: T_x M \rightarrow M$, is a mapping from the tangent space $T_x M$ to M with the requirement that a vector $\xi \in T_x M$ is mapped to the point $y := \exp_x(\xi) \in M$ such that there exists a geodesic $\gamma: [0, 1] \rightarrow M$, which satisfies $\gamma(0) = x$, $\gamma(1) = y$, and $\dot{\gamma}(0) = \xi$, where $\dot{\gamma}$ is the derivative of γ (see [14], [15]). Moreover, $\log_x: M \rightarrow T_x M$ denotes a logarithmic map at a point $x \in M$, which is defined as the inverse mapping of the exponential map at $x \in M$. For all $x, y \in M$, the existence of $\log_x(y)$ is guaranteed [16, Chapter V, Theorem 4.1] [17, Proposition 2.1].

Next, we give the definitions of a geodesically convex set and function (see [15, Section 2]) that generalize the concepts of a convex set and function in Euclidean space.

Definition II.1 (Geodesically convex set). *Let X be a subset of a Riemannian manifold M . X is said to be geodesically convex if, for any two points in X , there is a unique minimizing geodesic within X which joins those two points.*

Definition II.2 (Geodesically convex function). *A smooth function $f: M \rightarrow \mathbb{R}$ is said to be geodesically convex if, for any $x, y \in M$, it holds that*

$$f(y) \geq f(x) + \langle \text{grad } f(x), \log_x(y) \rangle_x,$$

where $\langle \cdot, \cdot \rangle_x$ is the Riemannian metric on M , and $\text{grad } f(x)$ is the Riemannian gradient of f at a point $x \in M$ (see [14]).

For $i \in \{1, 2, \dots, N\}$, let M_i be a Riemannian manifold and M be the Cartesian product of n Riemannian manifolds M_i (i.e., $M := M_1 \times \dots \times M_N$). $x^i \in M_i$ denotes a corresponding component of $x \in M$, and $\langle \cdot, \cdot \rangle_{x^i}$ denotes a Riemannian metric at a point $x^i \in M_i$. Furthermore, $\|\cdot\|_{x^i}$ represents the norm determined from the Riemann metric

$\langle \cdot, \cdot \rangle_{x^i}$. For a geodesically convex set $X_i \subset M_i$, we define the projection operator as $\Pi_{X_i}: M_i \rightarrow X_i$; i.e., $\Pi_{X_i}(x^i)$ is the unique point $y^i \in X_i$ minimizing $d^i(x^i, \cdot)$, where $d^i(\cdot, \cdot): M_i \times M_i \rightarrow \mathbb{R}$ denotes the distance function of M_i . The tangent space at a point $x = (x^1, x^2, \dots, x^N) \in M$ is given by $T_x M = T_{x^1} M_1 \oplus \dots \oplus T_{x^N} M_N$, by considering $T_{x^i} M_i$ to be a subspace of $T_x M$, where \oplus is the direct sum of vector spaces. Then, for a point $x = (x^1, x^2, \dots, x^N) \in M$ and a tangent vector $\xi \in T_x M$, we write $\xi = (\xi^i) = (\xi^1, \xi^2, \dots, \xi^N)$, where $i \in \{1, 2, \dots, N\}$, and $\xi^i \in T_{x^i} M_i$. Finally, for $x^i, y^i \in M_i$, $\varphi_{x^i \rightarrow y^i}^i$ denotes an isometry from $T_{x^i} M_i$ to $T_{y^i} M_i$ (e.g., $\varphi_{x^i \rightarrow y^i}^i$ stands for parallel transport from $T_{x^i} M_i$ to $T_{y^i} M_i$).

$\mathbb{E}[X]$ denotes the expectation of a random variable X , and $t_{[n]}$ denotes the history of the process up to time n (i.e., $t_{[n]} := (t_1, t_2, \dots, t_n)$). $\mathbb{E}[X|t_{[n]}]$ denotes the conditional expectation of X given $t_{[n]}$.

Assumption II.1. *For $i \in \{1, 2, \dots, N\}$, let M_i be a complete simply connected Riemannian manifold with sectional curvature lower bounded by $\kappa_i \leq 0$. We define $M := M_1 \times \dots \times M_N$. Then, we assume*

- (A1) *For all $i \in \{1, 2, \dots, N\}$, let $X_i \subset M_i$ be a bounded, closed, geodesically convex set¹ and $X := X_1 \times \dots \times X_N$. In addition, $X_i \subset M_i$ has a diameter bounded by D ; i.e., there exists a positive real number D such that*

$$\max_{i \in \{1, 2, \dots, N\}} \sup \{d^i(x^i, y^i) : x^i, y^i \in X_i\} \leq D,$$

where $d^i(\cdot, \cdot)$ denotes the distance function of M_i ;

- (A2) *A smooth function $f_t: M \rightarrow \mathbb{R}$ is geodesically convex, where t is a random variable whose probability distribution is a uniform distribution and supported on a set $\mathcal{T} := \{1, 2, \dots, T\}$. The function f is defined for all $x \in M$, by $f(x) := \mathbb{E}[f_t(x)] = (1/T) \sum_{t=1}^T f_t(x)$.*

Note that, when we define a positive number G as

$$G := \sup_{t \in \mathcal{T}, x \in X} \|\text{grad } f_t(x)\|_x,$$

we find that $G < \infty$ from Assumption II.1 (A1). The following is the main problem considered here [13, Section 4]:

Problem II.1. *Suppose that Assumption II.1 holds. Then, we have*

$$x_* \in X_* := \left\{ x_* \in X : f(x_*) = \inf_{x \in X} f(x) \right\}.$$

B. Background and motivation

Euclidean adaptive optimization algorithms, such as AdaGrad, Adam, and AMSGrad, are powerful tools for training deep neural networks. However, there are optimization problems on Riemannian manifolds in machine learning [4] that cannot be solved by Euclidean adaptive optimization algorithms. Accordingly, useful algorithms, such as RSGD [5,

¹The closedness and geodesical convexity of X_i imply the uniqueness and existence of $\Pi_{X_i}(x^i)$ [17, Proposition 2.4] [18, Theorem 1].

Section 2], RAdaGrad [13, Section 3.2], RAdam [13, Section 4], and RAMSGrad [13, Figure 1(a)], have been developed to solve Riemannian optimization problems. The algorithms, such as RAdaGrad, RAdam, and RAMSGrad, are based on the Euclidean adaptive optimization algorithms, AdaGrad, Adam, and AMSGrad. Hence, we can expect that the corresponding Riemannian adaptive optimization algorithms perform better than RSGD, the most basic Riemannian stochastic optimization algorithm. In fact, the numerical comparisons in [13] showed that Riemannian adaptive optimization algorithms are superior for the task of embedding the WordNet taxonomy in the Poincaré ball.

Although Riemannian adaptive optimization algorithms have been shown to be useful for Riemannian optimization in machine learning, we have two motivations related to the previous results in [13]. The first motivation is to identify whether or not RAMSGrad can solve directly the Riemannian optimization problem. This is because the previous results in [13] only showed that RAMSGrad performs a regret minimization, which does not lead to Riemannian minimization (see Subsection III-B for details). The second motivation is to identify whether or not RAMSGrad is applicable to significant problems in fields such as natural language processing and principal component analysis. This is because the previous results in [13] only gave a convergence analysis of RAMSGrad with a diminishing learning rate. Diminishing learning rates are approximately zero after a large number of iterations, which implies that algorithms with them are not implementable in practice. In contrast, a constant learning rate does not cause this problem (see Subsection III-B for details).

Therefore, our goal is to devise an algorithm that overcomes the above issues. In particular, the following section proposes an adaptive optimization algorithm (Algorithm 1) with a constant learning rate that can solve Problem II.1 directly (Corollary III.1).

III. RIEMANNIAN ADAPTIVE OPTIMIZATION ALGORITHM

A. Proposed algorithm and its convergence analysis

We propose the following algorithm (Algorithm 1). A small constant $\epsilon > 0$ in the definition of \hat{v}_n^i guarantees that $\sqrt{\hat{v}_n^i} > 0$ (Adam [8, Algorithm 1] and AMSGrad [10, Algorithm 2], [11, Algorithm 1] use such a constant in practice).

Now, let us compare Algorithm 1 on a Riemannian manifold with AMSGrad in Euclidean space. For simplicity, let us suppose that $M_i = X_i = \mathbb{R}$ ($i = 1, 2, \dots, N$). Then, Algorithm 1 defined on $M = \mathbb{R}^N$ is as follows: given $\mathbf{x}_1 \in \mathbb{R}^N$ and $\mathbf{m}_0 = \mathbf{v}_0 = \hat{\mathbf{v}}_0 = \mathbf{0} \in \mathbb{R}^N$,

$$\begin{aligned} \mathbf{m}_n &= \beta_{1n} \mathbf{m}_{n-1} + (1 - \beta_{1n}) \mathbf{g}_{t_n}, \\ \mathbf{v}_n &= \beta_2 \mathbf{v}_{n-1} + (1 - \beta_2) \mathbf{g}_{t_n} \odot \mathbf{g}_{t_n}, \\ \hat{\mathbf{v}}_n &= (\hat{v}_n^i)_i = (\max\{\hat{v}_{n-1}^i, v_n^i\} + \epsilon)_i, \\ \mathbf{x}_{n+1} &= \left(x_n^i - \alpha_n \frac{m_n^i}{\sqrt{\hat{v}_n^i}} \right)_i, \end{aligned}$$

where $\mathbf{x} \odot \mathbf{x} := (x^i)^2_i$ for $\mathbf{x} = (x^i)_i \in \mathbb{R}^N$. This implies Algorithm 1 is an extension of AMSGrad.

Algorithm 1 Modified RAMSGrad for solving Problem II.1

Require: $(\alpha_n)_{n \in \mathbb{N}} \subset [0, 1), (\beta_{1n})_{n \in \mathbb{N}} \subset [0, 1), \beta_2 \in [0, 1), \epsilon > 0$
1: $n \leftarrow 1, \mathbf{x}_1 \in X, \tau_0 = \mathbf{m}_0 = \mathbf{0} \in T_{x_0} M, v_0^i, \hat{v}_0^i = 0 \in \mathbb{R}$
2: **loop**
3: $\mathbf{g}_{t_n} = (g_{t_n}^i) = \text{grad } f_{t_n}(x_n)$
4: **for** $i = 1, 2, \dots, N$ **do**
5: $m_n^i = \beta_{1n} \tau_{n-1}^i + (1 - \beta_{1n}) g_{t_n}^i$
6: $v_n^i = \beta_2 v_{n-1}^i + (1 - \beta_2) \|g_{t_n}^i\|_{x_n^i}^2$
7: $\hat{v}_n^i = \max\{\hat{v}_{n-1}^i, v_n^i\} + \epsilon$
8: $x_{n+1}^i = \Pi_{X^i} \left[\exp_{x_n^i}^i \left(-\alpha_n \frac{m_n^i}{\sqrt{\hat{v}_n^i}} \right) \right]$
9: $\tau_n^i = \varphi_{x_n^i \rightarrow x_{n+1}^i}^i(m_n^i)$
10: **end for**
11: $n \leftarrow n + 1$
12: **end loop**

Our convergence analysis (Theorem III.1) allows Algorithm 1 to use both constant and diminishing learning rates. Corollaries III.1, and III.2 are convergence analyses of Algorithm 1 with constant and diminishing learning rates, respectively.

Theorem III.1. *Suppose that Assumption II.1 holds. Let $(x_n)_{n \in \mathbb{N}}$ and $(\hat{v}_n)_{n \in \mathbb{N}}$ be the sequences generated by Algorithm 1. We assume $\beta_{1n} \leq \beta_{1, n-1}$ for all $n \in \mathbb{N}$, and $(\alpha_n)_{n \in \mathbb{N}}$ is a sequence of positive learning rates, which satisfies $\alpha_n(1 - \beta_{1n}) \leq \alpha_{n-1}(1 - \beta_{1, n-1})$ for all $n \in \mathbb{N}$. We define $G := \max_{t \in \mathcal{T}, x \in X} \|\text{grad } f_t(x)\|_x$. Then, for all $x_* \in X_*$,*

$$\begin{aligned} & \mathbb{E} \left[\frac{1}{n} \sum_{k=1}^n f(x_k) - f(x_*) \right] \\ & \leq \frac{NGD^2}{2(1 - \beta_{11}) n \alpha_n} + \frac{G^2}{2\sqrt{\epsilon}(1 - \beta_{11})} \sum_{i=1}^N \zeta(\kappa_i, D) \frac{1}{n} \sum_{k=1}^n \alpha_k \\ & \quad + \frac{NGD}{1 - \beta_{11}} \frac{1}{n} \sum_{k=1}^n \beta_{1k}, \end{aligned} \tag{1}$$

where $\zeta(\kappa_i, D)$ is defined as in Lemma A.1.

Proof. See Appendix B. \square

Corollary III.1 (Constant learning rate). *Suppose that the assumptions in Theorem III.1 hold, $\alpha_n := \alpha > 0$, and $\beta_{1n} := \beta \in [0, 1)$. Then, Algorithm 1 satisfies, for all $x_* \in X_*$,*

$$\mathbb{E} \left[\frac{1}{n} \sum_{k=1}^n f(x_k) - f(x_*) \right] \leq \mathcal{O} \left(\frac{1}{n} \right) + C_1 \alpha + C_2 \beta,$$

where $C_1, C_2 > 0$ are constants.

Proof. See Appendix C. \square

Corollary III.2 (Diminishing learning rate). *Suppose that the assumptions in Theorem III.1 hold, $\alpha_n := 1/n^\eta$, where $\eta \in$*

$[1/2, 1)$, and $\sum_{k=1}^{\infty} \beta_{1k} < \infty$ ². Then, Algorithm 1 satisfies, for all $x_* \in X_*$,

$$\mathbb{E} \left[\frac{1}{n} \sum_{k=1}^n f(x_k) - f(x_*) \right] = \mathcal{O} \left(\frac{1}{n^{1-\eta}} \right).$$

Proof. See Appendix C. \square

B. Comparison of Algorithm 1 with the existing algorithms

Algorithm 1 with $n = t \in \mathcal{T}$ coincides with RAMSGrad [13, Figure 1(a)]. In [13], Bécigneul and Ganea used “regret” to guarantee the convergence of RAMSGrad. The regret at the end of T iterations is defined as

$$R_T := \sum_{t \in \mathcal{T}} f_t(x_t) - f_*,$$

where $(f_t)_{t \in \mathcal{T}}$ is a family of differentiable, geodesically convex functions from M to \mathbb{R} , $f_* := \min_{x \in X} \sum_{t \in \mathcal{T}} f_t(x)$, and $(x_t)_{t \in \mathcal{T}}$ is the sequence generated by RAMSGrad. They proved the following theorem [13, Theorem 1]:

Theorem III.2 (Convergence of RAMSGrad). *Suppose that Assumption II.1 (A1) holds and that f_t is smooth and geodesically convex for all $t \in \mathcal{T}$. Let $(x_t)_{t \in \mathcal{T}}$ and $(\hat{v}_t)_{t \in \mathcal{T}}$ be the sequences obtained from RAMSGrad, $\alpha_t = \alpha/\sqrt{t}$, $\beta_1 = \beta_{11}$, $\beta_{1k} \leq \beta_1$ for all $t \in \mathcal{T}$, $\alpha > 0$, and $\gamma := \beta_1/\sqrt{\beta_2} < 1$. We then have:*

$$R_T \leq \frac{\sqrt{T}D^2}{2\alpha(1-\beta_1)} \sum_{i=1}^N \sqrt{\hat{v}_T^i} + \frac{D^2}{2(1-\beta_1)} \sum_{i=1}^N \sum_{t=1}^T \beta_{1t} \frac{\sqrt{\hat{v}_t^i}}{\alpha_t} \\ + \frac{\alpha\sqrt{1+\log T}}{(1-\beta_1)^2(1-\gamma)\sqrt{1-\beta_2}} \sum_{i=1}^N \frac{\zeta(\kappa_i, D) + 1}{2} \sqrt{\sum_{t=1}^T \|g_{x_t}^i\|^2}.$$

Note that Theorem III.2 asserts the regret generated by RAMSGrad has an upper bound. We should also note that regret minimization does not always lead to solutions of Problem II.1. This is because, even if $(x_t)_{t \in \mathcal{T}}$ satisfies, for a sufficiently large number T ,

$$R_T = \sum_{t \in \mathcal{T}} f_t(x_t) - f_* \approx 0,$$

and we do *not* have that

$$Tf(x_T) - f_* = \sum_{t \in \mathcal{T}} f_t(x_T) - f_* \approx 0.$$

Accordingly, Theorem III.2 does not guarantee that the output x_T generated by RAMSGrad approximates the solution of Problem II.1. Additionally, Theorem III.2 assumes a diminishing learning rate α_t and does not assert anything about a constant learning rate.

Meanwhile, Corollary III.1 implies that, if we use sufficiently small constant learning rates α and β , then Algorithm 1 satisfies

$$\mathbb{E} \left[\frac{1}{n} \sum_{k=1}^n f(x_k) - f_* \right] \leq \mathcal{O} \left(\frac{1}{n} \right) + C_1\alpha + C_2\beta \approx \mathcal{O} \left(\frac{1}{n} \right),$$

² $\alpha_n := 1/n^\eta$ ($\eta \in [1/2, 1)$), and $\beta_{1n} = \lambda^n$ ($\lambda \in [0, 1)$) satisfy $\sum_{k=1}^{\infty} \beta_{1k} < \infty$, $\beta_{1n} \leq \beta_{1, n-1}$, and $\alpha_n(1-\beta_{1n}) \leq \alpha_{n-1}(1-\beta_{1, n-1})$ ($n \in \mathbb{N}$).

which implies that Algorithm 1 approximates the solution of Problem II.1 in the sense of the mean value of $f(x_k)$. Although Theorem III.2 can only use diminishing learning rates such that $\alpha_t := \alpha/\sqrt{t}$, Corollary III.1 guarantees that Algorithm 1 with a constant learning rate can solve Problem II.1.

Corollary III.2 implies that Algorithm 1 with a diminishing learning rate can solve Problem II.1 in the sense that

$$\mathbb{E} \left[\frac{1}{n} \sum_{k=1}^n f(x_k) - f_* \right] = \mathcal{O} \left(\frac{1}{n^{1-\eta}} \right),$$

while Theorem III.2 implies that RAMSGrad only minimizes the regret in the sense of the existence of a positive real number C such that

$$\frac{R_T}{T} \leq C\sqrt{\frac{1+\log T}{T}}.$$

Additionally, Theorem III.2 implies RAMSGrad only works in the case where $\eta = 1/2$, but Corollary III.2 implies Algorithm 1 works for a wider range of η .

The advantage of Corollary III.2 over Corollary III.1 is that using a diminishing learning rate is a robust way to solve Problem II.1. However, it is possible that Algorithm 1 with a diminishing learning rate does not work for a sufficiently large number S of iterations, because step 8 in Algorithm 1 with $\alpha_S \approx 0$ satisfies

$$x_{S+1}^i = \Pi_{X_i} \left[\exp_{x_S^i}^i \left(-\alpha_S \frac{m_S^i}{\sqrt{\hat{v}_S^i}} \right) \right] \approx x_S^i.$$

Such a trend was observed in [11]. The numerical results in [11] showed that Euclidean adaptive optimization algorithms, such as Adam and AMSGrad, with constant learning rates (e.g., $\alpha_n = \beta_{1n} = 10^{-3}$) perform better than those with diminishing ones in terms of both the training loss and accuracy score. Moreover, we can see that the Euclidean adaptive optimization algorithms in `torch.optim`³ use constant learning rates, such as $\alpha_n = 10^{-3}$ and $\beta_{1n} = 0.9, 0.999$.

According to [5, Section 2] and [13, Section 5], useful constant learning rates in RAMSGrad are $\alpha_n = 0.3, 0.1$ and $\beta_{1n} = 0.9$. Meanwhile, Corollary III.1 indicates that using a small constant learning rate β_{1n} would be good for solving Problem II.1. Accordingly, the next section numerically compares the behavior of Algorithm 1 with $\beta_{1n} = 0.9$ with one with $\beta_{1n} = 0.001 < 0.9$. Corollary III.2 (see also Theorem III.2) indicates that Algorithm 1 should use diminishing learning rates such that $\alpha_n = \mathcal{O}(1/\sqrt{n})$ and $\beta_{1n} = \lambda^n$ ($\lambda \in [0, 1)$). The next section uses diminishing learning rates to compare fairly the behaviors of Algorithm 1 with constant learning rates (see Section IV for details).

IV. NUMERICAL EXPERIMENTS

We numerically compared the following Riemannian stochastic optimization algorithms: RSGD [5, Section 2], RAdaGrad [13, Section 3.2], RAdam [13, Section 4], and Algorithm 1 (modified RAMSGrad). RAdam is obtained by removing the max operation in Algorithm 1, i.e., replacing

³<https://pytorch.org/docs/stable/optim.html>

$\hat{v}_n^i = \max\{\hat{v}_{n-1}^i, v_n^i\} + \epsilon$ with $\hat{v}_n^i = v_n^i + \epsilon$ (see [13]). Our experiments were conducted on a fast scalar computation server⁴ at Meiji University. The environment has two Intel(R) Xeon(R) Gold 6148 (2.4 GHz, 20 cores) CPUs, an NVIDIA Tesla V100 (16GB, 900Gbps) GPU and a Red Hat Enterprise Linux 7.6 operating system.

A. Poincaré embeddings

In [4], Nickel and Kiela developed Poincaré embeddings. Before describing the numerical experiments, we will review the fundamentals of hyperbolic geometry (see [4], [13], [19], [20]). $\mathcal{B}^d := \{x \in \mathbb{R}^d : \|x\| < 1\}$ denotes the open d -dimensional unit ball, where $\|\cdot\|$ denotes the Euclidean norm. The Poincaré ball model of hyperbolic space (\mathcal{B}^d, ρ) is defined by a manifold \mathcal{B}^d equipped with the following Riemannian metric:

$$\rho_x := \frac{4}{(1 - \|x\|^2)^2} \rho_x^E,$$

where $x \in \mathcal{B}^d$, and ρ_x^E denotes the Euclidean metric tensor. The Riemannian manifold (\mathcal{B}^d, ρ) has a constant sectional curvature, -1 . We define Möbius addition [19, Definition 1.10] of x and y in \mathcal{B}^d as

$$x \oplus_M y := \frac{(1 + 2\langle x, y \rangle + \|y\|^2)x + (1 - \|x\|^2)y}{1 + 2\langle x, y \rangle + \|x\|^2\|y\|^2},$$

where $\langle \cdot, \cdot \rangle := \rho^E(\cdot, \cdot)$. Moreover, $\ominus_M x$ denotes the left inverse [19, Definition 1.7] of $x \in \mathcal{B}^d$, and the Möbius gyrations [19, Definition 1.11] of \mathcal{B}^d are defined as

$$\text{gyr}[x, y]z := \ominus_M(x \oplus_M y) \oplus_M \{x \oplus_M (y \oplus_M z)\},$$

for all $x, y, z \in \mathcal{B}^d$.

In accordance with the above statements, the induced distance function on (\mathcal{B}^d, ρ) (see [20, Eq. (6)]) is defined for all $x, y \in \mathcal{B}^d$, by

$$d(x, y) = 2 \tanh^{-1}(\|(-x) \oplus_M y\|). \quad (2)$$

The exponential map on (\mathcal{B}^d, ρ) (see [20, Lemma 2]) is expressed as follows: for $x \in \mathcal{B}^d$ and $\xi \neq 0 \in T_x \mathcal{B}^d$,

$$\exp_x(\xi) = x \oplus_M \left\{ \tanh \left(\frac{\|\xi\|}{1 - \|x\|^2} \right) \right\} \frac{\xi}{\|\xi\|},$$

and, for $x \in \mathcal{B}^d$ and $0 \in T_x \mathcal{B}^d$,

$$\exp_x(0) = x.$$

Parallel transport of (\mathcal{B}^d, ρ) (see [13, Section 5]) along the unique geodesic from x to y is given by

$$\varphi_{x \rightarrow y}(\xi) = \frac{1 - \|y\|^2}{1 - \|x\|^2} \text{gyr}[y, -x]\xi.$$

The Riemannian gradient on (\mathcal{B}^d, ρ) (see [13, Section 5]) is expressed in terms of rescaled Euclidean gradients, i.e., for $x \in \mathcal{B}^d$, and the smooth function $f : \mathcal{B}^d \rightarrow \mathbb{R}$,

$$\text{grad } f(x) = \frac{(1 - \|x\|^2)^2}{4} \nabla^E f(x),$$

where $\nabla^E f(x)$ denotes the Euclidean gradient of f .

To compute the Poincaré embeddings for a set of N symbols by finding the embeddings $\Theta = \{u_i\}_{i=1}^N$, where $u_i \in \mathcal{B}^d$, we solve the following optimization problem: given $\mathcal{L} : \mathcal{B}^d \times \dots \times \mathcal{B}^d \rightarrow \mathbb{R}$,

$$\text{minimize } \mathcal{L}(\Theta) \quad \text{subject to } u_i \in \mathcal{B}^d. \quad (3)$$

The transitive closure of the WordNet mammals subtree consists of 1,180 nouns and 6,450 hypernymy Is-A relations. Let $\mathcal{D} = \{(u, v)\}$ be the set of observed hypernymy relations between noun pairs. We minimize a loss function defined by

$$\mathcal{L}(\Theta) = \sum_{(u,v) \in \mathcal{D}} \log \frac{e^{-d(u,v)}}{\sum_{v' \in \mathcal{N}(u)} e^{-d(u,v')}} \quad (4)$$

where $d(u, v)$ defined by (2) is the corresponding distance of the relation $(u, v) \in \mathcal{D}$, and $\mathcal{N}(u) = \{v' : (u, v') \notin \mathcal{D}\} \cup \{v\}$ is the set of negative examples for u including v (see [4], [13]). We embed the transitive closure of the WordNet mammals subtree into a 5-dimensional Poincaré ball (\mathcal{B}^5, ρ) .

Let us define $M_i := \mathcal{B}^5$ and $X_i := \{x \in \mathcal{B}^5 : \|x\| \leq 1 - 10^{-5}\}$, whose projection operator $\Pi_{X_i} : \mathcal{B}^5 \rightarrow X_i$ is computed as

$$\Pi_{X_i}(x) := \begin{cases} x & \text{if } \|x\| \leq 1 - 10^{-5} \\ (1 - 10^{-5}) \frac{x}{\|x\|} & \text{otherwise} \end{cases}.$$

Moreover, the geodesically convex set X_i has a bounded diameter; in fact, let D be the diameter of a closed disk X_i , measured by the Riemann metric of ρ .

As in [4], we will introduce an index for evaluating the embedding. For each observed relation $(u, v) \in \mathcal{D}$, we compute the corresponding distance $d(u, v)$ in the embedding and rank it among the set of negative relations for u , i.e., among the set $\{d(u, v') : (u, v') \notin \mathcal{D}\}$. In addition, we assume the reconstruction setting (see [4]); i.e., we evaluate the ranking of all nouns in the dataset. Then, we record the mean rank of v as well as the mean average precision (MAP) of the ranking. Thus, we evaluate the embedding in terms of the loss function values and the MAP rank.

We experimented with a special iteration called the ‘‘burn-in phase’’ (see [4, Section 3]) for the first 20 epochs. During the burn-in phase, the algorithm runs at a reduced learning rate of 1/100. When we minimized the loss function (4), we randomly sampled 10 negative relations per positive relation. We set $\epsilon = 10^{-8}$ in Algorithm 1.

The experiment used the code of Facebook Research⁵, and we used the NumPy 1.17.3 package and PyTorch 1.3.0 package.

1) *Constant learning rate*: First, we compared algorithms with the following ten constant learning rates:

(CS1) RSGD: $\alpha_n = 0.3$.

(CS2) RSGD: $\alpha_n = 0.1$.

(CG1) RAdaGrad: $\alpha_n = 0.3$.

(CG2) RAdaGrad: $\alpha_n = 0.1$.

(CD1) RAdam: $\alpha_n = 0.3, \beta_{1n} = 0.9, \beta_2 = 0.999$.

(CD2) RAdam: $\alpha_n = 0.1, \beta_{1n} = 0.9, \beta_2 = 0.999$.

⁴<https://www.meiji.ac.jp/isys/hpc/ia.html>

⁵<https://github.com/facebookresearch/poincare-embeddings>

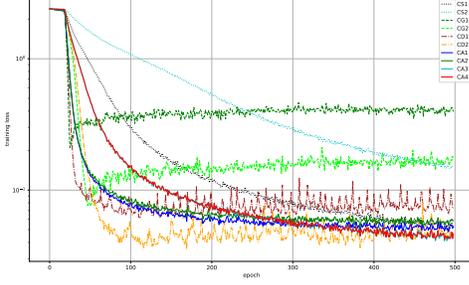


Fig. 1. Loss function value versus number of epochs in the case of constant learning rates.

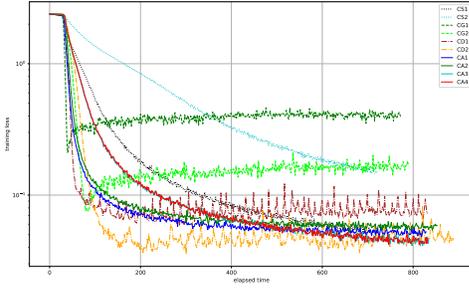


Fig. 2. Loss function value versus elapsed time in the case of constant learning rates.

- (CA1) Algorithm 1: $\alpha_n = 0.3$, $\beta_{1n} = 0.9$, $\beta_2 = 0.999$.
 (CA2) Algorithm 1: $\alpha_n = 0.3$, $\beta_{1n} = 0.001$, $\beta_2 = 0.999$.
 (CA3) Algorithm 1: $\alpha_n = 0.1$, $\beta_{1n} = 0.9$, $\beta_2 = 0.999$.
 (CA4) Algorithm 1: $\alpha_n = 0.1$, $\beta_{1n} = 0.001$, $\beta_2 = 0.999$.

The parameter α_n in (CS1) and (CS2) represents the learning rate of RSGD [5, Section 2]. The learning rates of (CA1)–(CA4) satisfy the assumptions of Corollary III.1. The parameters $\beta_2 = 0.999$ and $\beta_{1n} = 0.9$ in (CA1) and (CA3) are used in [13, Section 5]. We used $\beta_{1n} = 0.001$ in (CA2) and (CA4) to compare (CA1) and (CA3) with Algorithm 1 with a small learning rate. Figs. 1–4 show the numerical results. Fig. 1 shows the performances of the algorithms for loss function values defined by (4) with respect to the number of epochs, while Fig. 2 presents those with respect to the elapsed time. Fig. 3 shows the MAP ranks of the embeddings with respect to the number of epochs, while Fig. 4 presents the MAP ranks with respect to the elapsed time. We can see that Algorithm 1 outperforms RSGD and RAdaGrad in every setting. In particular, Figs. 1–2 show that the learning outcomes of RSGD fluctuate greatly depending on the learning rate. In contrast, Algorithm 1 and RAdam eventually reduce the loss function the most for any learning rate. Moreover, these figures show that the performance of (CA1) (resp. (CA3)) is comparable to that of (CA2) (resp. (CA4)). Meanwhile, RAdaGrad quickly reduced the objective function value in the early stages; however, it soon stopped learning.

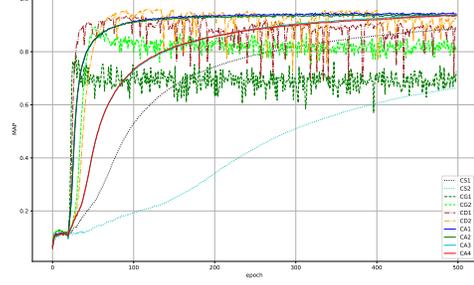


Fig. 3. MAP rank versus number of epochs in the case of constant learning rates.

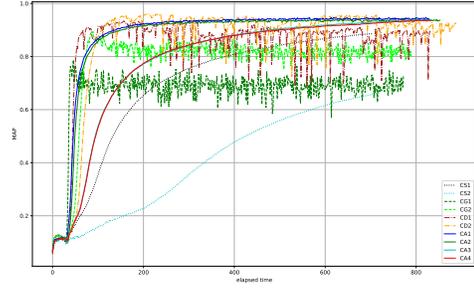


Fig. 4. MAP rank versus elapsed time in the case of constant learning rates.

2) *Diminishing learning rate*: Next, we compared algorithms with the following ten diminishing learning rates:

- (DS1) RSGD: $\alpha_n = 30/\sqrt{n}$.
 (DS2) RSGD: $\alpha_n = 10/\sqrt{n}$.
 (DG1) RAdaGrad: $\alpha_n = 30/\sqrt{n}$.
 (DG2) RAdaGrad: $\alpha_n = 10/\sqrt{n}$.
 (DD1) RAdam: $\alpha_n = 30/\sqrt{n}$, $\beta_{1n} = 0.5^n$, $\beta_2 = 0.999$.
 (DD2) RAdam: $\alpha_n = 10/\sqrt{n}$, $\beta_{1n} = 0.5^n$, $\beta_2 = 0.999$.
 (DA1) Algorithm 1: $\alpha_n = 30/\sqrt{n}$, $\beta_{1n} = 0.5^n$, $\beta_2 = 0.999$.
 (DA2) Algorithm 1: $\alpha_n = 30/\sqrt{n}$, $\beta_{1n} = 0.9^n$, $\beta_2 = 0.999$.
 (DA3) Algorithm 1: $\alpha_n = 10/\sqrt{n}$, $\beta_{1n} = 0.5^n$, $\beta_2 = 0.999$.
 (DA4) Algorithm 1: $\alpha_n = 10/\sqrt{n}$, $\beta_{1n} = 0.9^n$, $\beta_2 = 0.999$.

The learning rates of (DA1)–(DA4) satisfy the assumptions of Corollary III.2. We implemented (DA2) and (DA4) to compare them with (CA1) and (CA3). We implemented (DA1) and (DA3) to check how well Algorithm 1 works depending on the choice of β_{1n} . Figs. 5–8 show the numerical results. Fig. 5 shows the behaviors of the algorithms for loss function values defined by (4) with respect to the number of epochs, whereas Fig. 6 shows those with respect to the elapsed time. Fig. 7 presents the MAP ranks of the embeddings with respect to the number of epochs, while Fig. 8 shows MAP ranks with respect to the elapsed time. Even in the case of diminishing learning rates, Algorithm 1 outperforms RSGD in every setting. The learning results of RSGD fluctuate greatly depending on the initial learning rate. In particular, (DS2) reduces the loss function more slowly than the other algorithms do. On the other hand, Algorithm 1 and RAdam stably reduce the loss

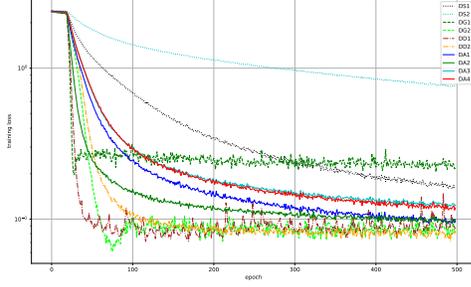


Fig. 5. Loss function value versus number of epochs in the case of diminishing learning rates.

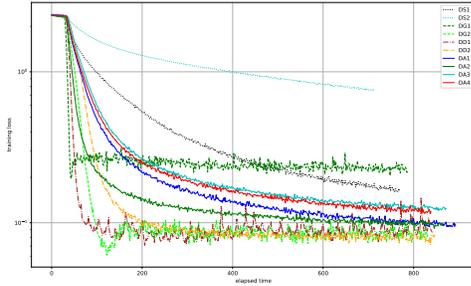


Fig. 6. Loss function value versus elapsed time in the case of diminishing learning rates.

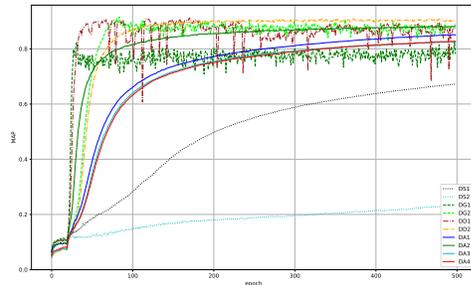


Fig. 7. MAP rank versus number of epochs in the case of diminishing learning rates.

function, regardless of the initial learning rate. Moreover, these figures indicate that (DA2) outperforms (DA1) and that (DA3) performs comparably to (DA4). In addition, RAdaGrad is better or worse than Algorithm 1 depending on how we choose the initial learning rates.

From Figs. 2 and 6, we can see that (CA1) (resp. (CA3)) outperforms (DA2) (resp. (DA4)). The above discussion shows that Algorithm 1 with a constant learning rate is superior to the other algorithms at embedding the WordNet mammals subtree into a Poincaré ball.

B. Principal component analysis

Here, we applied the algorithms to a principal component analysis (PCA) problem. Given n data points $a_1, \dots, a_n \in$

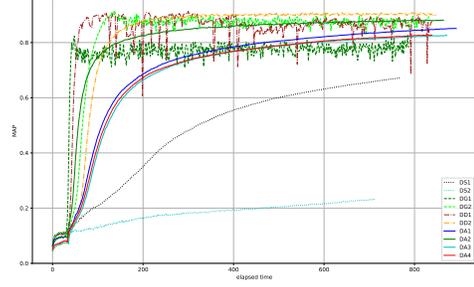


Fig. 8. MAP rank versus elapsed time in the case of diminishing learning rates.

\mathbb{R}^d , the PCA problem (see [12], [21]) is formulated as

$$\text{minimize } f(U) \quad \text{subject to } U \in \text{St}(k, d),$$

where

$$f(U) := -\frac{1}{n} \sum_{i=1}^n a_i^\top U U^\top a_i,$$

and $\text{St}(k, d) := \{U \in \mathbb{R}^{d \times k} : U^\top U = I_k\}$ denotes the Stiefel manifold. For this problem, we set $N = 1$ and $M = \text{St}(k, d)$. Since it is known that parallel transport has no closed-form solution on the Stiefel manifold, we use QR-based retraction and the associated vector transport as an approximation of the exponential map and the parallel transport, respectively (see [14]). The QR-based retraction is defined as

$$R_U(\xi) := \text{qf}(U + \xi),$$

where $U \in \text{St}(k, d)$, $\xi \in T_U \text{St}(k, d)$ and $\text{qf}(A)$ denotes the Q factor of the QR decomposition of A . Then, the associated vector transport is defined as

$$\mathcal{T}_{U \rightarrow V}(\xi) := \xi - V \text{sym}(V^\top \xi),$$

where $U, V \in \text{St}(k, d)$, $\xi \in T_U \text{St}(k, d)$ and $\text{sym}(A) := (A + A^\top)/2$. For this problem, the columns of the optimal solution U_* are known to be the top k eigenvectors of the data covariance matrix, which can be estimated using singular value decomposition. The performance of each algorithm in the experiment was judged in terms of the ‘‘optimality gap’’, that is, $f(U) - f(U_*)$. We evaluated the algorithms on the MNIST⁶ and digits⁷ datasets. The MNIST dataset contains handwritten digits data of 0–9 and has 10000 images of size 28×28 for testing (see [22]). For the MNIST dataset, we set $(n, k, d) = (10000, 784, 10)$. The digits dataset is made up of 1797 8×8 handwritten digit images. For the digits dataset, we set $(n, k, d) = (1797, 64, 8)$.

1) *Constant learning rate*: First, we compared ten algorithms (CS1)–(CA4) with constant learning rates, same as those used in the experiments on the Poincaré embeddings. Figs. 9–10 show the numerical results on the MNIST dataset, while Figs. 11–12 show the numerical results on the digits

⁶<https://keras.io/ja/datasets/>

⁷https://scikit-learn.org/stable/auto_examples/datasets/plot_digits_last_image.html

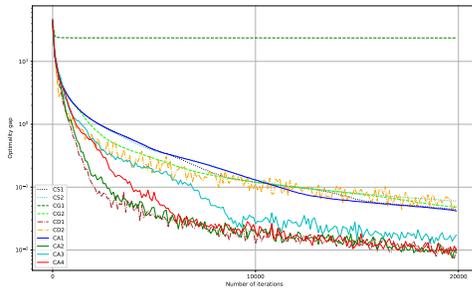


Fig. 9. Optimality gap versus number of iterations in the case of constant learning rates for the MNIST dataset

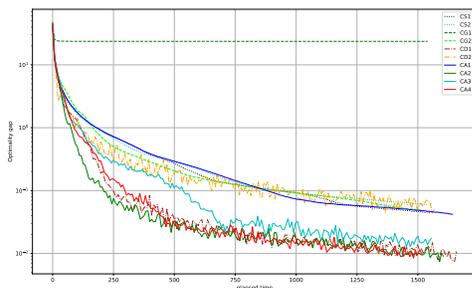


Fig. 10. Optimality gap versus elapsed time in the case of constant learning rates for the MNIST dataset

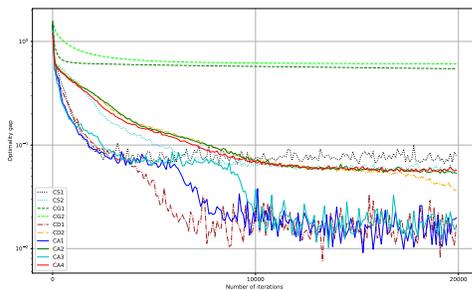


Fig. 11. Optimality gap versus number of iterations in the case of constant learning rates for the digits dataset

dataset. Figs. 9–10 indicate that Algorithm 1 performed well in every setting for the MNIST dataset. In particular, the behaviors of (CA2) and (CA4) are the best of all algorithms and the behavior of (CD1) is comparable to them. Moreover, Figs. 11–12 show that Algorithm 1 outperforms RSGD and RAdaGrad in every setting for the digits dataset. These figures indicate that (CD1), (CA1) and (CA3) eventually made the optimal gap the smallest. Meanwhile, Figs. 9–12 show that RAdaGrad often failed to reduce the optimal gap.

Moreover, we examined the supervised learning performance. The classification model in this case was the Linear Support Vector Machine⁸ (Linear SVM) provided by the

⁸<https://scikit-learn.org/stable/modules/generated/sklearn.svm.SVC.html>

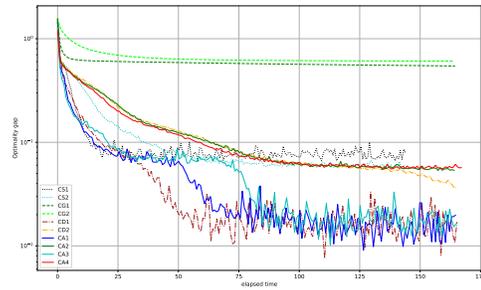


Fig. 12. Optimality gap versus elapsed time in the case of constant learning rates for the digits dataset

scikit-learn 0.23.2 package. TABLE I shows the 5-fold cross validation scores of the MNIST and digits dimensionally reduced by the last point generated by each algorithm with the constant learning rates. This table indicates that the algorithm which sufficiently minimizes the optimality gap also has high classification accuracy. For the MNIST dataset, since (CG1) does not converge to the optimal solution, its classification accuracy is also bad. Similarly, for the digits dataset, since (CG1) and (CG2) do not minimize the optimality gap, their classification accuracies are also bad.

TABLE I
THE CROSS VALIDATION SCORES OF THE LINEAR SVM IN THE CASE OF CONSTANT LEARNING RATES

	MNIST	digits
CS1	0.8109	0.8658
CS2	0.8104	0.8720
CG1	0.5992	0.7929
CG2	0.8164	0.7022
CD1	0.7973	0.8736
CD2	0.8078	0.8764
CA1	0.8168	0.8764
CA2	0.8099	0.8664
CA3	0.7931	0.8520
CA4	0.8131	0.8698

2) *Diminishing learning rate*: Next, we compared ten algorithms (DS1)–(DA4) with diminishing learning rates, same as those used in the experiments of the Poincaré embeddings. Figs. 13–14 show the numerical results on the MNIST dataset, while Figs. 15–16 show the numerical results on the digits dataset. Figs. 13–14 indicate that Algorithm 1 outperforms RSGD and RAdaGrad in every setting for the MNIST dataset. In particular, the behavior of (DD1) is the best of all and (DA1) performs comparably to (DD1). Moreover, Figs. 15–16 also show that Algorithm 1 outperforms RSGD and RAdaGrad in every setting for the digits dataset. In particular, the behaviors of (DA2) and (DA4) are the best of all algorithms and the behavior of (DD2) is comparable to them. Meanwhile, Figs. 13–16 show that RSGD and RAdaGrad often failed to reduce the optimal gap. Figs. 9 and 13 indicate that Algorithm 1 and RAdam with a constant learning rate are superior to the other algorithms for the MNIST dataset. Meanwhile, Figs. 11 and 15 indicate that Algorithm 1 and RAdam with a diminishing learning rate are superior to the other algorithms

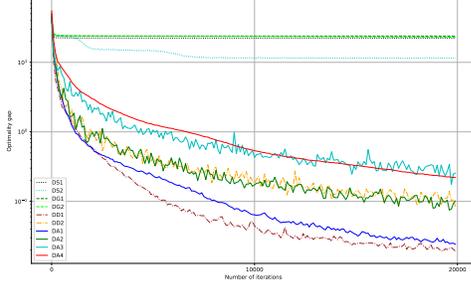


Fig. 13. Optimality gap versus number of iterations in the case of diminishing learning rates for the MNIST dataset

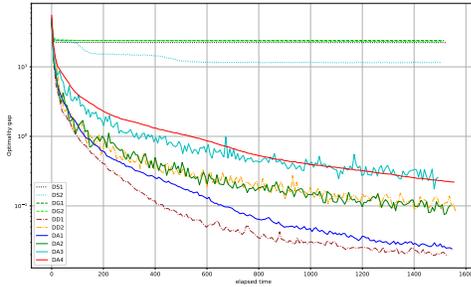


Fig. 14. Optimality gap versus elapsed time in the case of diminishing learning rates for the MNIST dataset

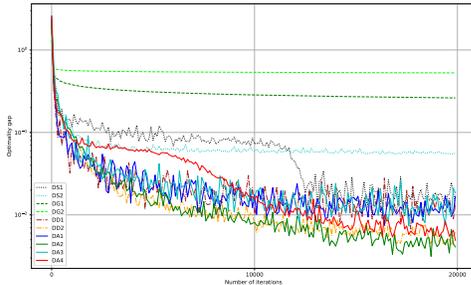


Fig. 15. Optimality gap versus number of iterations in the case of diminishing learning rates for the digits dataset

for the digits dataset.

As with the constant learning rate, we examined the supervised learning performance of the Linear SVM. TABLE II shows the 5-fold cross validation scores of the MNIST and digits dimensionally reduced by the last point generated by each algorithm with the diminishing learning rates. This table indicates that the algorithm which sufficiently minimizes the optimality gap also has high classification accuracy. Since, (DS1), (DS2), (DG1), and (DG2) for the MNIST dataset, and (DG1) and (DG2) for the digits dataset do not minimize the optimality gap, their classification accuracies are also bad.

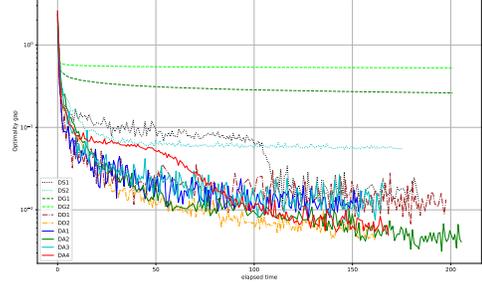


Fig. 16. Optimality gap versus elapsed time in the case of diminishing learning rates for the digits dataset

TABLE II
THE CROSS VALIDATION SCORES OF THE LINEAR SVM IN THE CASE OF DIMINISHING LEARNING RATES

	MNIST	digits
DS1	0.5719	0.8709
DS2	0.6332	0.8714
DG1	0.6546	0.7145
DG2	0.5904	0.6856
DD1	0.7955	0.8670
DD2	0.7928	0.8692
DA1	0.8133	0.8759
DA2	0.7922	0.8764
DA3	0.8239	0.8842
DA4	0.8061	0.8742

V. CONCLUSION

This paper proposed modified RAMSGrad, a Riemannian adaptive optimization method, and presented its convergence analysis. The proposed algorithm solves the Riemannian optimization problem directly, and it can use both constant and diminishing learning rates. We applied it to Poincaré embeddings and a PCA problem. The numerical experiments showed that it converges to the optimal solution faster than RSGD and RAdaGrad, and it minimizes the objective function regardless of the initial learning rate. In particular, an experiment showed that the proposed algorithm with a constant learning rate is a good way of embedding the WordNet mammals subtree into a Poincaré subtree. Moreover, we showed that, in the PCA problem, the choice between using a constant or a diminishing learning rate depends on the dataset.

VI. ACKNOWLEDGMENT

We are sincerely grateful to the editor and the anonymous referees for helping us improve the original manuscript.

APPENDIX A LEMMAS

Zhang and Sra developed the following lemma in [15, Lemma 5].

Lemma A.1 (Cosine inequality in Alexandrov spaces). *Let a, b, c be the sides (i.e., side lengths) of a geodesic triangle in an Alexandrov space whose curvature is bounded by $\kappa < 0$ and A be the angle between sides b and c . Then,*

$$a^2 \leq \zeta(\kappa, c)b^2 + c^2 - 2bc \cos(A),$$

where

$$\zeta(\kappa, c) = \frac{\sqrt{|\kappa|c}}{\tanh(\sqrt{|\kappa|c})}.$$

We will prove the following lemma. All relations between random variables hold almost surely.

Lemma A.2. *Suppose that Assumption II.1 (A2) holds. We define $G := \max_{t \in \mathcal{T}, x \in X} \|\text{grad } f_t(x)\|_x$. Let $(x_n)_{n \in \mathbb{N}}$ and $(\hat{v}_n)_{n \in \mathbb{N}}$ be the sequences generated by Algorithm 1. Then, for all $i \in \{1, 2, \dots, N\}$, and $k \in \mathbb{N}$,*

$$\|m_k^i\|_{x_k^i}^2 \leq G^2, \quad (5)$$

and

$$\sqrt{\hat{v}_k^i} \leq G. \quad (6)$$

Proof. First, we consider (5). The proof is by induction. For $k = 1$, from the convexity of $\|\cdot\|_{x_1^i}^2$, we have

$$\begin{aligned} \|m_1^i\|_{x_1^i}^2 &\leq \left\| \beta_{11} \varphi_{x_0^i \rightarrow x_1^i}^i(m_0^i) + (1 - \beta_{11})g_{t_1}^i \right\|_{x_1^i}^2 \\ &\leq \beta_{11} \left\| \varphi_{x_0^i \rightarrow x_1^i}^i(m_0^i) \right\|_{x_1^i}^2 + (1 - \beta_{11}) \|g_{t_1}^i\|_{x_1^i}^2 \\ &= (1 - \beta_{11}) \|g_{t_1}^i\|_{x_1^i}^2 \\ &\leq \|g_{t_1}^i\|_{x_1^i}^2 \\ &\leq G^2, \end{aligned}$$

where we have used $0 \leq \beta_{11} < 1$ and $\|g_{t_1}^i\|_{x_1^i} \leq G$. Suppose that $\|m_{k-1}^i\|_{x_{k-1}^i}^2 \leq G^2$. The convexity of $\|\cdot\|_{x_k^i}^2$, together with the definition of m_k^i , and $\|g_{t_k}^i\|_{x_k^i} \leq G$, guarantees that,

$$\begin{aligned} \|m_k^i\|_{x_k^i}^2 &\leq \beta_{1k} \left\| \varphi_{x_{k-1}^i \rightarrow x_k^i}^i(m_{k-1}^i) \right\|_{x_k^i}^2 + (1 - \beta_{1k}) \|g_{t_k}^i\|_{x_k^i}^2 \\ &\leq \beta_{1k} \|m_{k-1}^i\|_{x_{k-1}^i}^2 + (1 - \beta_{1k})G^2 \\ &\leq \beta_{1k}G^2 + (1 - \beta_{1k})G^2 \\ &= G^2. \end{aligned}$$

Thus, induction ensures that, for all $k \in \mathbb{N}$,

$$\|m_k^i\|_{x_k^i}^2 \leq G^2.$$

(6) can be proven in same way as (5). \square

APPENDIX B PROOF OF THEOREM III.1

Proof of Theorem III.1. Note that

$$y_{k+1}^i := \exp_{x_k^i}^i \left(-\alpha_k \frac{m_k^i}{\sqrt{\hat{v}_k^i}} \right).$$

Thus, we will consider a geodesic triangle consisting of three points x_k^i , x_*^i , and y_{k+1}^i . Let the length of each side be a , b , and c , respectively, such that

$$\begin{cases} a := d^i(y_{k+1}^i, x_*^i) \\ b := d^i(y_{k+1}^i, x_k^i) \\ c := d^i(x_k^i, x_*^i) \end{cases}. \quad (7)$$

It follows that

$$\begin{aligned} \cos(\angle y_{k+1}^i x_k^i x_*^i) &:= \frac{\left\langle \log_{x_k^i}^i(y_{k+1}^i), \log_{x_k^i}^i(x_*^i) \right\rangle_{x_k^i}}{\left\| \log_{x_k^i}^i(y_{k+1}^i) \right\|_{x_k^i} \left\| \log_{x_k^i}^i(x_*^i) \right\|_{x_k^i}} \\ &= \frac{\left\langle -\alpha_k \frac{m_k^i}{\sqrt{\hat{v}_k^i}}, \log_{x_k^i}^i(x_*^i) \right\rangle_{x_k^i}}{d^i(y_{k+1}^i, x_k^i) d^i(x_k^i, x_*^i)}. \end{aligned}$$

Using Lemma A.1 with (7) and the definition of Π_{X_i} , we have

$$\begin{aligned} &d^i(x_{k+1}^i, x_*^i)^2 \\ &\leq d^i(y_{k+1}^i, x_*^i)^2 \\ &\leq \zeta(\kappa^i, d^i(x_k^i, x_*^i)) d^i(y_{k+1}^i, x_k^i)^2 + d^i(x_k^i, x_*^i)^2 \\ &\quad - 2d^i(y_{k+1}^i, x_k^i) d^i(x_k^i, x_*^i) \frac{\left\langle -\alpha_k \frac{m_k^i}{\sqrt{\hat{v}_k^i}}, \log_{x_k^i}^i(x_*^i) \right\rangle_{x_k^i}}{d^i(y_{k+1}^i, x_k^i) d^i(x_k^i, x_*^i)}, \end{aligned}$$

which, together with the definition of y_{k+1}^i , implies that

$$\begin{aligned} \left\langle -m_k^i, \log_{x_k^i}^i(x_*^i) \right\rangle_{x_k^i} &\leq \frac{\sqrt{\hat{v}_k^i}}{2\alpha_k} (d^i(x_k^i, x_*^i)^2 - d^i(x_{k+1}^i, x_*^i)^2) \\ &\quad + \zeta(\kappa^i, d^i(x_k^i, x_*^i)) \frac{\alpha_k}{2\sqrt{\hat{v}_k^i}} \|m_k^i\|_{x_k^i}^2. \end{aligned}$$

Plugging $m_k^i = \beta_{1k} \varphi_{x_{k-1}^i \rightarrow x_k^i}^i(m_{k-1}^i) + (1 - \beta_{1k})g_{t_k}^i$ into the above inequality and using (A1), we obtain

$$\begin{aligned} &\left\langle -g_{t_k}^i, \log_{x_k^i}^i(x_*^i) \right\rangle_{x_k^i} \\ &\leq \frac{\sqrt{\hat{v}_k^i}}{2\alpha_k(1 - \beta_{1k})} (d^i(x_k^i, x_*^i)^2 - d^i(x_{k+1}^i, x_*^i)^2) \\ &\quad + \frac{\zeta(\kappa^i, D)}{2(1 - \beta_{1k})} \frac{\alpha_k}{\sqrt{\hat{v}_k^i}} \|m_k^i\|_{x_k^i}^2 \\ &\quad + \frac{\beta_{1k}}{1 - \beta_{1k}} \left\langle \varphi_{x_{k-1}^i \rightarrow x_k^i}^i(m_{k-1}^i), \log_{x_k^i}^i(x_*^i) \right\rangle_{x_k^i}. \end{aligned} \quad (8)$$

Since (A2) implies that f is geodesically convex with $g(x) = (g^i(x^i)) := \text{grad } f(x)$, we have

$$\begin{aligned} f(x_k) - f(x_*) &\leq \langle -g(x_k), \log_{x_k}^i(x_*) \rangle_{x_k} \\ &= \sum_{i=1}^N \left\langle -g^i(x_k^i), \log_{x_k^i}^i(x_*^i) \right\rangle_{x_k^i}. \end{aligned}$$

Summing the above equality from $k = 1$ to n , we obtain

$$\frac{1}{n} \sum_{k=1}^n f(x_k) - f(x_*) \leq \frac{1}{n} \sum_{k=1}^n \sum_{i=1}^N \left\langle -g^i(x_k^i), \log_{x_k^i}^i(x_*^i) \right\rangle_{x_k^i}. \quad (9)$$

Furthermore, the linearity of the Riemannian gradient ensures that

$$\begin{aligned} & \mathbb{E} \left[\left\langle -g_{t_k}^i, \log_{x_k^i}(x_*^i) \right\rangle_{x_k^i} \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\left\langle -g_{t_k}^i, \log_{x_k^i}(x_*^i) \right\rangle_{x_k^i} \mid t_{[k-1]} \right] \right] \\ &= \mathbb{E} \left[\left\langle -\mathbb{E} [g_{t_k}^i \mid t_{[k-1]}], \log_{x_k^i}(x_*^i) \right\rangle_{x_k^i} \right] \\ &= \mathbb{E} \left[\left\langle -g^i(x_k^i), \log_{x_k^i}(x_*^i) \right\rangle_{x_k^i} \right], \end{aligned}$$

which, together with (9), implies that

$$\begin{aligned} & \mathbb{E} \left[\frac{1}{n} \sum_{k=1}^n f(x_k) - f(x_*) \right] \\ & \leq \frac{1}{n} \mathbb{E} \left[\sum_{k=1}^n \sum_{i=1}^N \left\langle -g^i(x_k^i), \log_{x_k^i}(x_*^i) \right\rangle_{x_k^i} \right] \\ & = \frac{1}{n} \mathbb{E} \left[\sum_{k=1}^n \sum_{i=1}^N \left\langle -g_{t_k}^i, \log_{x_k^i}(x_*^i) \right\rangle_{x_k^i} \right]. \end{aligned}$$

From (8) and the above inequality, we have

$$\begin{aligned} & \mathbb{E} \left[\frac{1}{n} \sum_{k=1}^n f(x_k) - f(x_*) \right] \\ & \leq \frac{1}{n} \mathbb{E} \left[\sum_{k=1}^n \sum_{i=1}^N \frac{\sqrt{\hat{v}_k^i}}{2\alpha_k(1-\beta_{1k})} (d^i(x_k^i, x_*^i)^2 - d^i(x_{k+1}^i, x_*^i)^2) \right] \\ & \quad + \frac{1}{n} \mathbb{E} \left[\sum_{k=1}^n \sum_{i=1}^N \frac{\zeta(\kappa^i, D)}{2(1-\beta_{1k})} \frac{\alpha_k}{\sqrt{\hat{v}_k^i}} \|m_k^i\|_{x_k^i}^2 \right] \\ & \quad + \frac{1}{n} \mathbb{E} \left[\sum_{k=1}^n \sum_{i=1}^N \frac{\beta_{1k}}{1-\beta_{1k}} \left\langle \varphi_{x_{k-1}^i \rightarrow x_k^i}(m_{k-1}^i), \log_{x_k^i}(x_*^i) \right\rangle_{x_k^i} \right]. \end{aligned} \tag{10}$$

Here, let us consider the first term of the left-hand side of (10). We note that from the assumption for all $k \in \mathbb{N}$, $\alpha_k(1-\beta_{1k}) \leq \alpha_{k-1}(1-\beta_{1,k-1})$, and $\beta_{1k} \leq \beta_{1,k-1}$,

$$\alpha_k(1-\beta_{1k}) \leq \alpha_{k-1}(1-\beta_{1,k-1}) \leq \alpha_{k-1}(1-\beta_{1k}),$$

which implies $\alpha_k \leq \alpha_{k-1}$. Using $\beta_{1k} \leq \beta_{11}$, $\alpha_k \leq \alpha_{k-1}$, $\sqrt{\hat{v}_k^i} \geq \sqrt{\hat{v}_{k-1}^i}$, and $\alpha_k(1-\beta_{1k}) \leq \alpha_{k-1}(1-\beta_{1,k-1})$ for all $k \in \mathbb{N}$, together with (A1), we have that

$$\begin{aligned} & \sum_{k=1}^n \sum_{i=1}^N \frac{\sqrt{\hat{v}_k^i}}{2\alpha_k(1-\beta_{1k})} (d^i(x_k^i, x_*^i)^2 - d^i(x_{k+1}^i, x_*^i)^2) \\ & \leq \frac{1}{2(1-\beta_{11})} \sum_{i=1}^N \left[\sum_{k=2}^n \left(\frac{\sqrt{\hat{v}_k^i}}{\alpha_k} - \frac{\sqrt{\hat{v}_{k-1}^i}}{\alpha_{k-1}} \right) d^i(x_k^i, x_*^i)^2 \right. \\ & \quad \left. + \frac{\sqrt{\hat{v}_1^i}}{\alpha_1} d^i(x_1^i, x_*^i)^2 \right] \\ & \leq \frac{1}{2(1-\beta_{11})} \sum_{i=1}^N \left[\sum_{k=2}^n \left(\frac{\sqrt{\hat{v}_k^i}}{\alpha_k} - \frac{\sqrt{\hat{v}_{k-1}^i}}{\alpha_{k-1}} \right) D^2 + \frac{\sqrt{\hat{v}_1^i}}{\alpha_1} D^2 \right] \end{aligned}$$

$$\begin{aligned} & = \frac{D^2}{2(1-\beta_{11})} \sum_{i=1}^N \frac{\sqrt{\hat{v}_n^i}}{\alpha_n} \\ & \leq \frac{NGD^2}{2\alpha_n(1-\beta_{11})}, \end{aligned}$$

where the last inequality is guaranteed by Lemma A.2. Namely,

$$\begin{aligned} & \mathbb{E} \left[\sum_{k=1}^n \sum_{i=1}^N \frac{\sqrt{\hat{v}_k^i}}{2\alpha_k(1-\beta_{1k})} (d^i(x_k^i, x_*^i)^2 - d^i(x_{k+1}^i, x_*^i)^2) \right] \\ & \leq \frac{NGD^2}{2\alpha_n(1-\beta_{11})}. \end{aligned} \tag{11}$$

Next, let us consider the second term of the left-hand side of (10). From $\sqrt{\epsilon} \leq \sqrt{\hat{v}_k^i}$ and Lemma A.2, we have

$$\begin{aligned} & \sum_{k=1}^n \sum_{i=1}^N \frac{\zeta(\kappa^i, D)}{2(1-\beta_{1k})} \frac{\alpha_k}{\sqrt{\hat{v}_k^i}} \|m_k^i\|_{x_k^i}^2 \\ & \leq \frac{G^2}{2\sqrt{\epsilon}(1-\beta_{11})} \sum_{i=1}^N \zeta(\kappa_i, D) \sum_{k=1}^n \alpha_k. \end{aligned}$$

Namely,

$$\begin{aligned} & \mathbb{E} \left[\sum_{k=1}^n \sum_{i=1}^N \frac{\zeta(\kappa^i, D)}{2(1-\beta_{1k})} \frac{\alpha_k}{\sqrt{\hat{v}_k^i}} \|m_k^i\|_{x_k^i}^2 \right] \\ & \leq \frac{G^2}{2\sqrt{\epsilon}(1-\beta_{11})} \sum_{i=1}^N \zeta(\kappa_i, D) \sum_{k=1}^n \alpha_k. \end{aligned} \tag{12}$$

Now, let us consider the third term of the left-hand side of (10). Applying the Cauchy-Schwarz inequality to the term and using (A1) and Lemma A.2, it follows that

$$\begin{aligned} & \sum_{k=1}^n \sum_{i=1}^N \frac{\beta_{1k}}{1-\beta_{1k}} \left\langle \varphi_{x_{k-1}^i \rightarrow x_k^i}(m_{k-1}^i), \log_{x_k^i}(x_*^i) \right\rangle_{x_k^i} \\ & \leq \sum_{k=1}^n \sum_{i=1}^N \frac{\beta_{1k}}{1-\beta_{1k}} \left\| \varphi_{x_{k-1}^i \rightarrow x_k^i}(m_{k-1}^i) \right\|_{x_k^i} \left\| \log_{x_k^i}(x_*^i) \right\|_{x_k^i} \\ & \leq \frac{NGD}{1-\beta_{11}} \sum_{k=1}^n \beta_{1k}. \end{aligned}$$

Namely,

$$\begin{aligned} & \mathbb{E} \left[\sum_{k=1}^n \sum_{i=1}^N \frac{\beta_{1k}}{1-\beta_{1k}} \left\langle \varphi_{x_{k-1}^i \rightarrow x_k^i}(m_{k-1}^i), \log_{x_k^i}(x_*^i) \right\rangle_{x_k^i} \right] \\ & \leq \frac{NGD}{1-\beta_{11}} \sum_{k=1}^n \beta_{1k}. \end{aligned} \tag{13}$$

Finally, together with (10), (11), (12), and (13), we have

$$\begin{aligned} & \mathbb{E} \left[\frac{1}{n} \sum_{k=1}^n f(x_k) - f(x_*) \right] \\ & \leq \frac{NGD^2}{2(1-\beta_{11})} \frac{1}{n\alpha_n} + \frac{G^2}{2\sqrt{\epsilon}(1-\beta_{11})} \sum_{i=1}^N \zeta(\kappa_i, D) \frac{1}{n} \sum_{k=1}^n \alpha_k \\ & \quad + \frac{NGD}{1-\beta_{11}} \frac{1}{n} \sum_{k=1}^n \beta_{1k}. \end{aligned}$$

This complete the proof. \square

APPENDIX C
PROOF OF COROLLARY III.1 AND III.2

Proof of Corollary III.1. The learning rates $\alpha_n := \alpha$ and $\beta_{1n} := \beta$ satisfy for all $n \in \mathbb{N}$, $\beta_{1n} \leq \beta_{1,n-1}$ and $\alpha_n(1 - \beta_{1n}) \leq \alpha_{n-1}(1 - \beta_{1,n-1})$. Let us define

$$C_1 := \frac{G^2}{\sqrt{\epsilon}(1 - \beta_{11})} \sum_{i=1}^N \zeta(\kappa_i, D) > 0,$$

and

$$C_2 := \frac{NGD}{1 - \beta_{11}}.$$

Using the definitions of C_1 and C_2 , (1) can be written as

$$\mathbb{E} \left[\frac{1}{n} \sum_{k=1}^n f(x_k) - f(x_*) \right] \leq \frac{NGD^2}{2\alpha(1 - \beta_{11})} \frac{1}{n} + C_1\alpha + C_2\beta.$$

This complete the proof. \square

Proof of Corollary III.2. Let $\alpha_n = 1/n^\eta$ ($\eta \in [1/2, 1)$) and $(\beta_{1n})_{n \in \mathbb{N}}$ satisfies $\beta_{1n} \leq \beta_{1,n-1}$ and $\alpha_n(1 - \beta_{1n}) \leq \alpha_{n-1}(1 - \beta_{1,n-1})$ for all $n \in \mathbb{N}$, and $\sum_{k=1}^{\infty} \beta_{1k} < \infty$. First, we obviously have

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \beta_{1k} \leq \lim_{n \rightarrow \infty} \frac{B_1}{n} = 0, \quad (14)$$

where $B_1 := \sum_{k=1}^{\infty} \beta_{1k} < \infty$. We have that

$$\lim_{n \rightarrow \infty} \frac{1}{n\alpha_n} = \lim_{n \rightarrow \infty} \frac{1}{n^{1-\eta}} = 0.$$

Furthermore, we have

$$\frac{1}{n} \sum_{k=1}^n \alpha_k \leq \frac{1}{n} \left(1 + \int_1^n \frac{dt}{t^\eta} \right) \leq \frac{1}{1-\eta} \frac{1}{n^{1-\eta}}. \quad (15)$$

This, together with (1), (14), and (15), proves the assertion of Corollary III.2. \square

REFERENCES

- [1] A. Iosifidis, A. Tefas, and I. Pitas, "Graph embedded extreme learning machine," *IEEE Transactions on Cybernetics*, vol. 46, no. 1, pp. 311–324, 2015.
- [2] S. Mao, L. Xiong, L. Jiao, T. Feng, and S.-K. Yeung, "A novel Riemannian metric based on Riemannian structure and scaling information for fixed low-rank matrix completion," *IEEE Transactions on Cybernetics*, vol. 47, no. 5, pp. 1299–1312, 2016.
- [3] X. Shen and F.-L. Chung, "Deep network embedding for graph representation learning in signed networks," *IEEE Transactions on Cybernetics*, 2018.
- [4] M. Nickel and D. Kiela, "Poincaré embeddings for learning hierarchical representations," in *Advances in neural information processing systems*, 2017, pp. 6338–6347.
- [5] S. Bonnabel, "Stochastic gradient descent on Riemannian manifolds," *IEEE Transactions on Automatic Control*, vol. 58, no. 9, pp. 2217–2229, 2013.
- [6] H. Sato, H. Kasai, and B. Mishra, "Riemannian stochastic variance reduced gradient algorithm with retraction and vector transport," *SIAM Journal on Optimization*, vol. 29, no. 2, pp. 1444–1472, 2019.
- [7] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," *Journal of machine learning research*, pp. 2121–2159, 2011.

- [8] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *Proceedings of The International Conference on Learning Representations*, pp. 1–15, 2015.
- [9] M. D. Zeiler, "Adadelta: an adaptive learning rate method," *arXiv preprint arXiv:1212.5701*, 2012.
- [10] S. J. Reddi, S. Kale, and S. Kumar, "On the convergence of Adam and beyond," *Proceedings of The International Conference on Learning Representations*, pp. 1–23, 2018.
- [11] H. Iiduka, "Appropriate learning rates of adaptive learning rate optimization algorithms for training deep neural networks," *arXiv preprint arXiv:2002.09647*, 2020.
- [12] H. Kasai, P. Jawanpuria, and B. Mishra, "Riemannian adaptive stochastic gradient algorithms on matrix manifolds," in *International Conference on Machine Learning*, 2019, pp. 3262–3271.
- [13] G. Bécigneul and O.-E. Ganea, "Riemannian adaptive optimization methods," *Proceedings of The International Conference on Learning Representations*, 2019.
- [14] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization algorithms on matrix manifolds*. Princeton University Press, 2008.
- [15] H. Zhang and S. Sra, "First-order methods for geodesically convex optimization," in *Conference on Learning Theory*, 2016, pp. 1617–1638.
- [16] T. Sakai, *Riemannian geometry*. American Mathematical Society, 1996, vol. 149.
- [17] C. Li, G. López, and V. Martquez, "Iterative algorithms for nonexpansive mappings on Hadamard manifolds," *Taiwanese Journal of Mathematics*, vol. 14, no. 2, pp. 541–559, 2010.
- [18] R. Walter, "On the metric projection onto convex sets in riemannian spaces," *Archiv der Mathematik*, vol. 25, no. 1, pp. 91–98, 1974.
- [19] A. A. Ungar, "A gyrovector space approach to hyperbolic geometry," *Synthesis Lectures on Mathematics and Statistics*, vol. 1, no. 1, pp. 1–194, 2008.
- [20] O.-E. Ganea, G. Bécigneul, and T. Hofmann, "Hyperbolic neural networks," in *Advances in neural information processing systems*, 2018, pp. 5345–5355.
- [21] P. Zhou, X. Yuan, S. Yan, and J. Feng, "Faster first-order methods for stochastic non-convex optimization on Riemannian manifolds," *IEEE transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [22] Y. LeCun, C. Cortes, and C. J. Burges, "The mnist database," URL <http://yann.lecun.com/exdb/mnist>, 1998.



Hiroyuki Sakai has been an undergraduate student in the Department of Computer Science, School of Science and Technology, Meiji University, Kanagawa, Japan, since April 2017. His research field is optimization theory and its applications.



Hideaki Iiduka received the Ph.D. degree in mathematical and computing science from Tokyo Institute of Technology, Tokyo, Japan, in 2005. From 2005 to 2007, he was a Research Assistant in the Department of Mathematical and Computing Sciences, Tokyo Institute of Technology, Tokyo, Japan. From 2007 to 2008, he was a Research Fellow (PD) of the Japan Society for the Promotion of Science. From October 2008 to March 2013, he was an Associate Professor in the Network Design Research Center, Kyushu Institute of Technology, Tokyo, Japan. From April 2013 to March 2019, he was an Associate Professor in the Department of Computer Science, School of Science and Technology, Meiji University, Kanagawa, Japan. Since April 2019, he has been a Professor in the same department. His research field is optimization theory and its applications to mathematical information science. He is a member of SIAM and MOS.