
Escaping from saddle points on Riemannian manifolds

Yue Sun
University of Washington
Seattle, WA 98105
yuesun@uw.edu

Nicolas Flammarion
EPFL
Lausanne, Switzerland
nicolas.flammarion@epfl.ch

Maryam Fazel
University of Washington
Seattle, WA 98105
mfazel@uw.edu

Abstract

We consider minimizing a nonconvex, smooth function f on a Riemannian manifold \mathcal{M} . We show that a perturbed version of Riemannian gradient descent algorithm converges to a second-order stationary point (and hence is able to *escape* saddle points on the manifold). The rate of convergence depends as $1/\epsilon^2$ on the accuracy ϵ , which matches a rate known only for unconstrained smooth minimization. The convergence rate depends polylogarithmically on the manifold dimension d , hence is almost dimension-free. The rate also has a polynomial dependence on the parameters describing the curvature of the manifold and the smoothness of the function. While the unconstrained problem (Euclidean setting) is well-studied, our result is the first to prove such a rate for nonconvex, manifold-constrained problems.

1 Introduction

We consider minimizing a non-convex smooth function on a smooth manifold \mathcal{M} ,

$$\text{minimize}_{x \in \mathcal{M}} f(x), \tag{1}$$

where \mathcal{M} is a d -dimensional smooth manifold¹, and f is twice differentiable, with a Hessian that is ρ -Lipschitz (assumptions are formalized in section 4). This framework includes a wide range of fundamental problems (often non-convex), such as PCA (Edelman et al., 1998), dictionary learning (Sun et al., 2017), low rank matrix completion (Boumal & Absil, 2011), and tensor factorization (Ishteva et al., 2011). Finding the global minimum to Eq. (1) is in general NP-hard; our goal is to find an approximate second order stationary point with first order optimization methods. We are interested in first-order methods because they are extremely prevalent in machine learning, partly because computing Hessians is often too costly. It is then important to understand how first-order methods fare when applied to nonconvex problems, and there has been a wave of recent interest on this topic since (Ge et al., 2015), as reviewed below.

In the Euclidean space, it is known that with random initialization, gradient descent avoids saddle points asymptotically (Pemantle, 1990; Lee et al., 2016). Lee et al. (2017) (section 5.5) show that this is also true on smooth manifolds, although the result is expressed in terms of nonstandard manifold smoothness measures. Also, importantly, this line of work does not give quantitative rates for the algorithm's behaviour near saddle points.

¹Here d is the dimension of the manifold itself; we do not consider \mathcal{M} as a submanifold of a higher dimensional space. For instance, if \mathcal{M} is a 2-dimensional sphere embedded in \mathbb{R}^3 , its dimension is $d = 2$.

Du et al. (2017) show gradient descent can be *exponentially slow* in the presence of saddle points. To alleviate this phenomenon, it is shown that for a β -gradient Lipschitz, ρ -Hessian Lipschitz function, cubic regularization (Carmon & Duchi, 2017) and perturbed gradient descent (Ge et al., 2015; Jin et al., 2017a) converges to $(\epsilon, -\sqrt{\rho\epsilon})$ local minimum² in polynomial time, and momentum based method accelerates (Jin et al., 2017b). Much less is known about inequality constraints: Nouiehed et al. (2018) and Mokhtari et al. (2018) discuss second order convergence for general inequality-constrained problems, where they need an NP-hard subproblem (checking the co-positivity of a matrix) to admit a polynomial time approximation algorithm. However such an approximation exists only under very restrictive assumptions.

An orthogonal line of work is optimization on Riemannian manifolds. Absil et al. (2009) provide comprehensive background, showing how algorithms such as gradient descent, Newton and trust region methods can be implemented on Riemannian manifolds, together with asymptotic convergence guarantees to first order stationary points. Zhang & Sra (2016) provide global convergence guarantees for first order methods when optimizing geodesically convex functions. Bonnabel (2013) obtains the first asymptotic convergence result for stochastic gradient descent in this setting, which is further extended by Tripuraneni et al. (2018); Zhang et al. (2016); Khuzani & Li (2017). If the problem is non-convex, or the Riemannian Hessian is not positive definite, one can use second order methods to escape from saddle points. Boumal et al. (2016a) shows that Riemannian trust region method converges to a second order stationary point in polynomial time (see, also, Kasai & Mishra, 2018; Hu et al., 2018; Zhang & Zhang, 2018). But this method requires a Hessian oracle, whose complexity is d times more than computing gradient. In Euclidean space, trust region subproblem can be sometimes solved via a Hessian-vector product oracle, whose complexity is about the same as computing gradients. Agarwal et al. (2018) discuss its implementation on Riemannian manifolds, but not clear about the complexity and sensitivity of Hessian vector product oracle on manifold.

The study of the convergence of gradient descent for non-convex Riemannian problems is previously done only in the Euclidean space by modeling the manifold with equality constraints. Ge et al. (2015, Appendix B) prove that stochastic projected gradient descent methods converge to second order stationary points in polynomial time (here the analysis is not geometric, and depends on the algebraic representation of the equality constraints). Sun & Fazel (2018) proves perturbed projected gradient descent converges with a comparable rate to the unconstrained setting (Jin et al., 2017a) (polylog in dimension). The paper applies projections from the ambient Euclidean space to the manifold and analyzes the iterations under the Euclidean metric. This approach loses the geometric perspective enabled by Riemannian optimization, and cannot explain convergence rates in terms of inherent quantities such as the sectional curvature of the manifold.

Contributions. We provide convergence guarantees for perturbed first order Riemannian optimization methods to second-order stationary points (local minima). We prove that as long as the function is appropriately smooth and the manifold has bounded sectional curvature, a perturbed Riemannian gradient descent algorithm escapes (an approximate) saddle points with a rate of $1/\epsilon^2$, a polylog dependence on the dimension of the manifold (hence almost dimension-free), and a polynomial dependence on the smoothness and curvature parameters. This is the first result showing such a rate for Riemannian optimization, and the first to relate the rate to geometric parameters of the manifold.

Despite analogies with the unconstrained (Euclidean) analysis and with the Riemannian optimization literature, the technical challenge in our proof goes beyond combining two lines of work: we need to analyze the interaction between the first-order method and the second order structure of the manifold to obtain second-order convergence guarantees that depend on the manifold curvature. Unlike in Euclidean space, the curvature affects the Taylor approximation of gradient steps. On the other hand, unlike in the local rate analysis in first-order Riemannian optimization, our second-order analysis requires more refined properties of the manifold structure (whereas in prior work, first order oracle makes enough progress for a local convergence rate proof, see Lemma 1), and second order algorithms such as (Boumal et al., 2016a) use second order oracles (Hessian evaluation). See section 4 for further discussion.

²defined as x satisfying $\|\nabla f(x)\| \leq \epsilon$, $\lambda_{\min} \nabla^2 f(x) \geq -\sqrt{\rho\epsilon}$

2 Notation and Background

We consider a complete³, smooth, d dimensional Riemannian manifold $(\mathcal{M}, \mathfrak{g})$, equipped with a Riemannian metric \mathfrak{g} , and we denote by $\mathcal{T}_x\mathcal{M}$ its tangent space at $x \in \mathcal{M}$ (which is a vector space of dimension d). We also denote by $\mathbb{B}_x(r) = \{v \in \mathcal{T}_x\mathcal{M}, \|v\| \leq r\}$ the ball of radius r in $\mathcal{T}_x\mathcal{M}$ centered at 0. At any point $x \in \mathcal{M}$, the metric \mathfrak{g} induces a natural inner product on the tangent space denoted by $\langle \cdot, \cdot \rangle : \mathcal{T}_x\mathcal{M} \times \mathcal{T}_x\mathcal{M} \rightarrow \mathbb{R}$. We also consider the Levi-Civita connection ∇ (Absil et al., 2009, Theorem 5.3.1). The Riemannian curvature tensor is denoted by $R(x)[u, v]$ where $x \in \mathcal{M}$, $u, v \in \mathcal{T}_x\mathcal{M}$ and is defined in terms of the connection ∇ (Absil et al., 2009, Theorem 5.3.1). The sectional curvature $K(x)[u, v]$ for $x \in \mathcal{M}$ and $u, v \in \mathcal{T}_x\mathcal{M}$ is then defined in Lee (1997, Prop. 8.8).

$$K(x)[u, v] = \frac{\langle R(x)[u, v]u, v \rangle}{\langle u, u \rangle \langle v, v \rangle - \langle u, v \rangle^2}, \quad x \in \mathcal{M}, \quad u, v \in \mathcal{T}_x\mathcal{M}.$$

Denote the distance (induced by the Riemannian metric) between two points in \mathcal{M} by $d(x, y)$. A geodesic $\gamma : \mathbb{R} \rightarrow \mathcal{M}$ is a constant speed curve whose length is equal to $d(x, y)$, so it is the shortest path on manifold linking x and y . $\gamma_{x \rightarrow y}$ denotes the geodesic from x to y (thus $\gamma_{x \rightarrow y}(0) = x$ and $\gamma_{x \rightarrow y}(1) = y$).

The exponential map $\text{Exp}_x(v)$ maps $v \in \mathcal{T}_x\mathcal{M}$ to $y \in \mathcal{M}$ such that there exists a geodesic γ with $\gamma(0) = x$, $\gamma(1) = y$ and $\frac{d}{dt}\gamma(0) = v$. The injectivity radius at point $x \in \mathcal{M}$ is the maximal radius r for which the exponential map is a diffeomorphism on $\mathbb{B}_x(r) \subset \mathcal{T}_x\mathcal{M}$. The injectivity radius of the manifold, denoted by \mathfrak{I} , is the infimum of the injectivity radii at all points. Since the manifold is complete, we have $\mathfrak{I} > 0$. When $x, y \in \mathcal{M}$ satisfies $d(x, y) \leq \mathfrak{I}$, the exponential map admits an inverse $\text{Exp}_x^{-1}(y)$, which satisfies $d(x, y) = \|\text{Exp}_x^{-1}(y)\|$. Parallel translation Γ_x^y denotes a the map which transports $v \in \mathcal{T}_x\mathcal{M}$ to $\Gamma_x^y v \in \mathcal{T}_y\mathcal{M}$ along $\gamma_{x \rightarrow y}$ such that the vector stays constant by satisfying a zero-acceleration condition (Lee, 1997, equation (4.13)).

For a smooth function $f : \mathcal{M} \rightarrow \mathbb{R}$, $\text{grad}f(x) \in \mathcal{T}_x\mathcal{M}$ denotes the Riemannian gradient of f at $x \in \mathcal{M}$ which satisfies $\frac{d}{dt}f(\gamma(t)) = \langle \gamma'(t), \text{grad}f(x) \rangle$ (see Absil et al., 2009, Sec 3.5.1 and (3.31)). The Hessian of f is defined jointly with the Riemannian structure of the manifold. The (directional) Hessian is $H(x)[\xi_x] := \nabla_{\xi_x} \text{grad}f$, and we use $H(x)[u, v] := \langle u, H(x)[v] \rangle$ as a shorthand. We call $x \in \mathcal{M}$ an $(\epsilon, -\sqrt{\rho\epsilon})$ saddle point when $\|\nabla f(x)\| \leq \epsilon$ and $\lambda_{\min}(H(x)) \leq -\sqrt{\rho\epsilon}$. We refer the interested reader to Do Carmo (2016) and Lee (1997) which provide a thorough review on these important concepts of Riemannian geometry.

3 Perturbed Riemannian gradient algorithm

Our main Algorithm 1 runs as follows:

1. Check the norm of the gradient: If it is large, do one step of Riemannian gradient descent, consequently the function value decreases.
2. If the norm of gradient is small, it's either an approximate saddle point or a local minimum. Perturb the variable by adding an appropriate level of noise in its tangent space, map it back to the manifold and run a few iterations.
 - (a) If the function value decreases, iterates are escaping from the approximate saddle point (and the algorithm continues)
 - (b) If the function value does not decrease, then it is an approximate local minimum (the algorithm terminates).

Algorithm 1 relies on the manifold's exponential map, and is useful for cases where this map is easy to compute⁴. We refer readers to Lee (1997, pp. 81-86) for the exponential map of sphere and hyperbolic manifolds, and Absil et al. (2009, Example 5.4.2, 5.4.3) for the Stiefel and Grassmann manifolds. If the exponential map is not computable, the algorithm can use a retraction⁵ instead,

³Since our results are local, completeness is not necessary and our results can be easily generalized, with extra assumptions on the injectivity radius.

⁴Numerous interesting manifolds have closed-form exponential maps: *the Grassmannian manifold, the Stiefel manifold, the Minkowski space, the hyperbolic space, SE(n), SO(n)*... (see, Miolane et al. (2018); Boumal et al. (2014) and their open-source packages and Bécigneul & Ganea (2019, Sec 5) for an example in NLP).

⁵A retraction is a first-order approximation of the exponential map which is often easier to compute.

Algorithm 1 Perturbed Riemannian gradient algorithm

Require: Initial point $x_0 \in \mathcal{M}$, parameters $\beta, \rho, K, \mathfrak{J}$, accuracy ϵ , probability of success δ (parameters defined in Assumptions 1, 2, 3 and assumption of Theorem 1).

Set constants: $\hat{c} \geq 4$, $C := C(K, \beta, \rho)$ (defined in Lemma 2 and proof of Lemma 8)

and $\sqrt{c_{\max}} \leq \frac{1}{56\hat{c}^2}$, $r = \frac{\sqrt{c_{\max}}}{\chi^2} \epsilon$, $\chi = 3 \max\{\log(\frac{d\beta(f(x_0) - f^*)}{\hat{c}\epsilon^2\delta}), 4\}$.

Set threshold values: $f_{\text{thres}} = \frac{c_{\max}}{\chi^3} \sqrt{\frac{\epsilon^3}{\rho}}$, $g_{\text{thres}} = \frac{\sqrt{c_{\max}}}{\chi^2} \epsilon$, $t_{\text{thres}} = \frac{\chi}{c_{\max}} \frac{\beta}{\sqrt{\rho\epsilon}}$, $t_{\text{noise}} = -t_{\text{thres}} - 1$.

Set stepsize: $\eta = \frac{c_{\max}}{\beta}$.

while 1 do

if $\|\text{grad}f(x_t)\| \leq g_{\text{thres}}$ **and** $t - t_{\text{noise}} > t_{\text{thres}}$ **then**

$t_{\text{noise}} \leftarrow t$, $\tilde{x}_t \leftarrow x_t$, $x_t \leftarrow \text{Exp}_{x_t}(\xi_t)$, ξ_t uniformly sampled from $\mathbb{B}_{x_t}(r) \subset \mathcal{T}_x\mathcal{M}$.

end if

if $t - t_{\text{noise}} = t_{\text{thres}}$ **and** $f(x_t) - f(\tilde{x}_{t_{\text{noise}}}) > -f_{\text{thres}}$ **then**

output $\tilde{x}_{t_{\text{noise}}}$

end if

$x_{t+1} \leftarrow \text{Exp}_{x_t}(-\min\{\eta, \frac{\mathfrak{J}}{\|\text{grad}f(x_t)\|}\}\text{grad}f(x_t))$.

$t \leftarrow t + 1$.

end while

however our current analysis only covers the case of the exponential map. In Figure 1, we illustrate a function with saddle point on sphere, and plot the trajectory of Algorithm 1 when it is initialized at a saddle point.




Figure 1: Function f with saddle point on a sphere. $f(x) = x_1^2 - x_2^2 + 4x_3^2$. We plot the contour of this function on unit sphere. Algorithm 1 initializes at $x_0 = [1, 0, 0]$ (a saddle point), perturbs it towards x_1 and runs Riemannian gradient descent, and terminates at $x^* = [0, -1, 0]$ (a local minimum). We amplify the first iteration to make saddle perturbation visible.

4 Main theorem: escape rate for perturbed Riemannian gradient descent

We now turn to our main results, beginning with our assumptions and a statement of our main theorem. We then develop a brief proof sketch.

Our main result involves two conditions on function f and one on the curvature of the manifold \mathcal{M} .

Assumption 1 (Lipschitz gradient). *There is a finite constant β such that*

$$\|\text{grad}f(y) - \Gamma_x^y \text{grad}f(x)\| \leq \beta d(x, y) \quad \text{for all } x, y \in \mathcal{M}.$$

Assumption 2 (Lipschitz Hessian). *There is a finite constant ρ such that*

$$\|H(y) - \Gamma_x^y H(x) \Gamma_y^x\|_2 \leq \rho d(x, y) \quad \text{for all } x, y \in \mathcal{M}.$$

Assumption 3 (Bounded sectional curvature). *There is a finite constant K such that*

$$|K(x)[u, v]| \leq K \quad \text{for all } x \in \mathcal{M} \text{ and } u, v \in \mathcal{T}_x\mathcal{M}$$

K is an intrinsic parameter of the manifold capturing the curvature. We list a few examples here: (i) A sphere of radius R has a constant sectional curvature $K = 1/R^2$ (Lee, 1997, Theorem 1.9). If

the radius is bigger, K is smaller which means the sphere is less curved; (ii) A hyper-bolic space H_R^n of radius R has $K = -1/R^2$ (Lee, 1997, Theorem 1.9); (iii) For sectional curvature of the Stiefel and the Grassmann manifolds, we refer readers to Rapcsák (2008, Section 5) and Wong (1968), respectively.

Note that the constant K is not directly related to the RLICQ parameter R defined by Ge et al. (2015) which first requires describing the manifold by equality constraints. Different representations of the same manifold could lead to different curvature bounds, while sectional curvature is an intrinsic property of manifold. If the manifold is a sphere $\sum_{i=1}^{d+1} x_i^2 = R^2$, then $K = 1/R^2$, but more generally there is no simple connection. The smoothness parameters we assume are natural compared to some quantity from complicated compositions Lee et al. (2017) (Section 5.5) or pullback (Zhang & Zhang, 2018). With these assumptions, the main result of this paper is the following:

Theorem 1. *Under Assumptions 1,2,3, let $C(K, \beta, \rho)$ be a function defined in Lemma 2, $\hat{\rho} = \max\{\rho, C(K, \beta, \rho)\}$, if ϵ satisfies that*

$$\epsilon \leq \min \left\{ \frac{\hat{\rho}}{56 \max\{c_2(K), c_3(K)\} \eta \beta} \log \left(\frac{d\beta}{\sqrt{\hat{\rho}\epsilon\delta}} \right), \left(\frac{\mathfrak{I}\hat{\rho}}{12\hat{c}\sqrt{\eta}\beta} \log \left(\frac{d\beta}{\sqrt{\hat{\rho}\epsilon\delta}} \right) \right)^2 \right\} \quad (2)$$

where $c_2(K), c_3(K)$ are defined in Lemma 4, then with probability $1 - \delta$, perturbed Riemannian gradient descent with step size c_{\max}/β converges to a $(\epsilon, -\sqrt{\hat{\rho}\epsilon})$ -stationary point of f in

$$O \left(\frac{\beta(f(x_0) - f(x^*))}{\epsilon^2} \log^4 \left(\frac{\beta d(f(x_0) - f(x^*))}{\epsilon^2 \delta} \right) \right)$$

iterations.

Proof roadmap. For a function satisfying smoothness condition (Assumption 1 and 2), we use a local upper bound of the objective based on the third-order Taylor expansion (see supplementary material Section A for a review),

$$f(u) \leq f(x) + \langle \text{grad}f(x), \text{Exp}_x^{-1}(u) \rangle + \frac{1}{2}H(x)[\text{Exp}_x^{-1}(u), \text{Exp}_x^{-1}(u)] + \frac{\rho}{6}\|\text{Exp}_x^{-1}(u)\|^3.$$

When the norm of the gradient is large (not near a saddle), the following lemma guarantees the decrease of the objective function in one iteration.

Lemma 1. (Boumal et al., 2018) *Under Assumption 1, by choosing $\bar{\eta} = \min\{\eta, \frac{\mathfrak{I}}{\|\text{grad}f(u)\|}\} = O(1/\beta)$, the Riemannian gradient descent algorithm is monotonically descending, $f(u^+) \leq f(u) - \frac{1}{2}\bar{\eta}\|\text{grad}f(u)\|^2$.*

Thus our main challenge in proving the main theorem is the Riemannian gradient behaviour at an approximate saddle point:

1. Similar to the Euclidean case studied by Jin et al. (2017a), we need to bound the ‘‘thickness’’ of the ‘‘stuck region’’ where the perturbation fails. We still use a pair of hypothetical auxiliary sequences and study the ‘‘coupling’’ sequences. When two perturbations couple in the thinnest direction of the stuck region, their distance grows and one of them escapes from saddle point.
2. However our iterates are evolving on a manifold rather than a Euclidean space, so our strategy is to map the iterates back to an appropriate fixed tangent space where we can use the Euclidean analysis. This is done using the inverse of the exponential map and various parallel transports.
3. Several key challenges arise in doing this. Unlike Jin et al. (2017a), the structure of the manifold interacts with the local approximation of the objective function in a complicated way. On the other hand, unlike recent work on Riemannian optimization by Boumal et al. (2016a), we do not have access to a second order oracle and we need to understand how the sectional curvature and the injectivity radius (which both capture intrinsic manifold properties) affect the behavior of the first order iterates.
4. Our main contribution is to carefully investigate how the various approximation errors arising from (a) the linearization of the iteration couplings and (b) their mappings to a common tangent space can be handled on manifolds with bounded sectional curvature. We address these challenges in a sequence of lemmas (Lemmas 3 through 6) we combine to linearize the coupling iterations in a common tangent space and precisely control the approximation error. This result is formally stated in the following lemma.

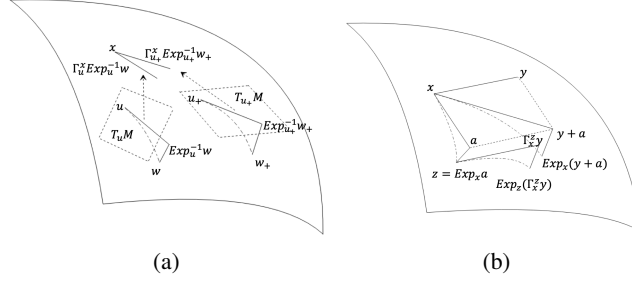


Figure 2: (a) Eq. (5). First map w and w_+ to $\mathcal{T}_u\mathcal{M}$ and $\mathcal{T}_{u_+}\mathcal{M}$, and transport the two vectors to $\mathcal{T}_x\mathcal{M}$, and get their relation. (b) Lemma 3 bounds the difference of two steps starting from x : (1) take $y + a$ step in $\mathcal{T}_x\mathcal{M}$ and map it to manifold, and (2) take a step in $\mathcal{T}_x\mathcal{M}$, map to manifold, call it z , and take $\Gamma_x^z y$ step in $\mathcal{T}_x\mathcal{M}$, and map to manifold. $\text{Exp}_z(\Gamma_x^z y)$ is close to $\text{Exp}_x(y + a)$.

Lemma 2. Define $\gamma = \sqrt{\hat{\rho}\epsilon}$, $\kappa = \frac{\beta}{\gamma}$, and $\mathcal{S} = \sqrt{\eta\beta}\frac{\gamma}{\hat{\rho}} \log^{-1}(\frac{d\kappa}{\delta})$. Let us consider x be a $(\epsilon, -\sqrt{\hat{\rho}\epsilon})$ saddle point, and define $u^+ = \text{Exp}_u(-\eta\text{grad}f(u))$ and $w^+ = \text{Exp}_w(-\eta\text{grad}f(w))$. Under Assumptions 1, 2, 3, if all pairwise distances between u, w, u^+, w^+, x are less than $12\mathcal{S}$, then for some explicit constant $C(K, \rho, \beta)$ depending only on K, ρ, β , there is

$$\begin{aligned} & \|\text{Exp}_x^{-1}(w^+) - \text{Exp}_x^{-1}(u^+) - (I - \eta H(x))(\text{Exp}_x^{-1}(w) - \text{Exp}_x^{-1}(u))\| \\ & \leq C(K, \rho, \beta)d(u, w) (d(u, w) + d(u, x) + d(w, x)). \end{aligned} \quad (3)$$

The proof of this lemma includes novel contributions by strengthen known result (Lemmas 3) and also combining known inequalities in novel ways (Lemmas 4 to 6) that allow us to control all the approximation errors and arrive at the tight rate of escape for the algorithm.

5 Proof of Lemma 2

Lemma 2 controls the error of the linear approximation of the iterates when mapped in $\mathcal{T}_x\mathcal{M}$. In this section, we assume that all points are within a region of diameter $R := 12\mathcal{S} \leq \mathfrak{J}$ (inequality follows from Eq. (2)), i.e., the distance of any two points in the following lemmas are less than R . The proof of Lemma 2 is based on the sequence of following lemmas.

Lemma 3. Let $x \in \mathcal{M}$ and $y, a \in \mathcal{T}_x\mathcal{M}$. Let us denote by $z = \text{Exp}_x(a)$ then under Assumption 3

$$d(\text{Exp}_x(y + a), \text{Exp}_z(\Gamma_x^z y)) \leq c_1(K) \min\{\|a\|, \|y\|\}(\|a\| + \|y\|)^2. \quad (4)$$

This lemma tightens the result of Karcher (1977, C2.3), which only shows an upper-bound $O(\|a\|(\|a\| + \|y\|)^2)$. We prove the upper-bound $O(\|y\|(\|a\| + \|y\|)^2)$ in the supplement. We also need the following lemma showing that both the exponential map and its inverse are Lipschitz.

Lemma 4. Let $x, y, z \in \mathcal{M}$, and the distance of each two points is no bigger than R . Then under assumption 3

$$(1 + c_2(K)R^2)^{-1}d(y, z) \leq \|\text{Exp}_x^{-1}(y) - \text{Exp}_x^{-1}(z)\| \leq (1 + c_3(K)R^2)d(y, z).$$

Intuitively this lemma relates the norm of the difference of two vectors of $\mathcal{T}_x\mathcal{M}$ to the distance between the corresponding points on the manifold \mathcal{M} and follows from bounds on the Hessian of the square-distance function (Sakai, 1996, Ex. 4 p. 154). The upper-bound is directly proven by Karcher (1977, Proof of Cor. 1.6), and we prove the lower-bound via Lemma 3 in the supplement.

The following contraction result is fairly classical and is proven using the Rauch comparison theorem from differential geometry (Cheeger & Ebin, 2008).

Lemma 5. (Mangoubi et al., 2018, Lemma 1) Under Assumption 3, for $x, y \in \mathcal{M}$ and $w \in \mathcal{T}_x\mathcal{M}$,

$$d(\text{Exp}_x(w), \text{Exp}_y(\Gamma_x^y w)) \leq c_4(K)d(x, y).$$

Finally we need the following corollary of the Ambrose-Singer theorem (Ambrose & Singer, 1953).

Lemma 6. (Karcher, 1977, Section 6) Under Assumption 3, for $x, y, z \in \mathcal{M}$ and $w \in \mathcal{T}_x\mathcal{M}$,

$$\|\Gamma_y^z \Gamma_x^y w - \Gamma_x^z w\| \leq c_5(K)d(x, y)d(y, z)\|w\|.$$

Lemma 3 through 6 are mainly proven in the literature, and we make up the missing part in Supplementary material Section B. Then we prove Lemma 2 in Supplementary material Section B.

The spirit of the proof is to linearize the manifold using the exponential map and its inverse, and to carefully bounds the various error terms caused by the approximation. Let us denote by $\theta = d(u, w) + d(u, x) + d(w, x)$.

1. We first show using twice Lemma 3 and Lemma 5 that

$$d(\text{Exp}_u(\text{Exp}_u^{-1}(w) - \eta \Gamma_w^u \text{grad} f(w)), \text{Exp}_u(-\eta \text{grad} f(u) + \Gamma_{u_+}^u \text{Exp}_{u_+}^{-1}(w_+))) = O(\theta d(u, w)).$$

2. We use Lemma 4 to linearize this iteration in $\mathcal{T}_u \mathcal{M}$ as

$$\|\Gamma_{u_+}^u \text{Exp}_{u_+}^{-1}(w_+) - \text{Exp}_u^{-1}(w) + \eta[\text{grad} f(u) - \Gamma_w^u \text{grad} f(w)]\| = O(\theta d(u, w)).$$

3. Using the Hessian Lipschitzness

$$\|\Gamma_{u_+}^u \text{Exp}_{u_+}^{-1}(w_+) - \text{Exp}_u^{-1}(w) + \eta H(u) \text{Exp}_u^{-1}(w)\| = O(\theta d(u, w)).$$

3. We use Lemma 6 to map to $T_x \mathcal{M}$ and the Hessian Lipschitzness to compare $H(u)$ to $H(x)$. This is an important intermediate result (see Lemma 1 in Supplementary material Section B).

$$\|\Gamma_{u_+}^x \text{Exp}_{u_+}^{-1}(w_+) - \Gamma_u^x \text{Exp}_u^{-1}(w) + \eta H(x) \Gamma_u^x \text{Exp}_u^{-1}(w)\| = O(\theta d(u, w)). \quad (5)$$

4. We use Lemma 3 and 4 to approximate two iteration updates in $\mathcal{T}_x \mathcal{M}$.

$$\|\text{Exp}_x^{-1}(w) - (\text{Exp}_x^{-1}(u) + \Gamma_u^x \text{Exp}_u^{-1}(w))\| \leq O(\theta d(u, w)). \quad (6)$$

And same for the u_+, w_+ pair replacing u, w .

5. Combining Eq. (5) and Eq. (6) together, we obtain

$$\|\text{Exp}_x^{-1}(w^+) - \text{Exp}_x^{-1}(u^+) - (I - \eta H(x))(\text{Exp}_x^{-1}(w) - \text{Exp}_x^{-1}(u))\| \leq O(\theta d(u, w)).$$

Now note that, the iterations u, u_+, w, w_+ of the algorithm are both on the manifold. We use $\text{Exp}_x^{-1}(\cdot)$ to map them to the same tangent space at x .

Therefore we have linearized the two coupled trajectories $\text{Exp}_x^{-1}(u_t)$ and $\text{Exp}_x^{-1}(w_t)$ in a common tangent space, and we can modify the Euclidean escaping saddle analysis thanks to the error bound we proved in Lemma 2.

6 Proof of main theorem

In this section we suppose all assumptions in Section 4 hold. The proof strategy is to show with high probability that the function value decreases of \mathcal{F} in \mathcal{T} iterations at an approximate saddle point. Lemma 7 suggests that, if after a perturbation and \mathcal{T} steps, the iterate is $\Omega(\mathcal{S})$ far from the approximate saddle point, then the function value decreases. If the iterates do not move far, the perturbation falls in a stuck region. Lemma 8 uses a coupling strategy, and suggests that the width of the stuck region is small in the negative eigenvector direction of the Riemannian Hessian.

Define

$$\mathcal{F} = \eta \beta \frac{\gamma^3}{\hat{\rho}^2} \log^{-3}\left(\frac{d\kappa}{\delta}\right), \quad \mathcal{G} = \sqrt{\eta \beta} \frac{\gamma^2}{\hat{\rho}} \log^{-2}\left(\frac{d\kappa}{\delta}\right), \quad \mathcal{T} = \frac{\log\left(\frac{d\kappa}{\delta}\right)}{\eta \gamma}.$$

At an approximate saddle point \tilde{x} , let y be in the neighborhood of \tilde{x} where $d(y, \tilde{x}) \leq \mathfrak{I}$, denote

$$\tilde{f}_y(x) := f(y) + \langle \text{grad} f(y), \text{Exp}_y^{-1}(\tilde{x}) \rangle + \frac{1}{2} \Gamma_{\tilde{x}}^y H(\tilde{x}) \Gamma_y^{\tilde{x}} [\text{Exp}_y^{-1}(\tilde{x}), \text{Exp}_y^{-1}(\tilde{x})].$$

Let $\|\text{grad} f(\tilde{x})\| \leq \mathcal{G}$ and $\lambda_{\min}(H(\tilde{x})) \leq -\gamma$. We consider two iterate sequences, u_0, u_1, \dots and w_0, w_1, \dots where u_0, w_0 are two perturbations at \tilde{x} .

Lemma 7. *Assume Assumptions 1, 2, 3 and Eq. (2) hold. There exists a constant c_{\max} , $\forall \hat{c} > 3, \delta \in (0, \frac{d\kappa}{\hat{c}}]$, for any u_0 with $d(\tilde{x}, u_0) \leq 2\mathcal{S}/(\kappa \log(\frac{d\kappa}{\delta}))$, $\kappa = \beta/\gamma$.*

$$T = \min \left\{ \inf_t \left\{ t | \tilde{f}_{u_0}(u_t) - f(u_0) \leq -3\mathcal{F} \right\}, \hat{c}\mathcal{T} \right\},$$

then $\forall \eta \leq c_{\max}/\beta$, we have $\forall 0 < t < T$, $d(u_0, u_t) \leq 3(\hat{c}\mathcal{S})$.

Lemma 8. Assume Assumptions 1, 2, 3 and Eq. (2) hold. Take two points u_0 and w_0 which are perturbed from an approximate saddle point, where $d(\tilde{x}, u_0) \leq 2\mathcal{L}/(\kappa \log(\frac{\kappa}{\delta}))$, $\text{Exp}_{\tilde{x}}^{-1}(w_0) - \text{Exp}_{\tilde{x}}^{-1}(u_0) = \mu r e_1$, e_1 is the smallest eigenvector⁶ of $H(\tilde{x})$, $\mu \in [\delta/(2\sqrt{d}), 1]$, and the algorithm runs two sequences $\{u_t\}$ and $\{w_t\}$ starting from u_0 and w_0 . Denote

$$T = \min \left\{ \inf_t \left\{ t | \tilde{f}_{w_0}(w_t) - f(w_0) \leq -3\mathcal{F} \right\}, \hat{c}\mathcal{T} \right\},$$

then $\forall \eta \leq c_{\max}/l$, if $\forall 0 < t < T$, $d(\tilde{x}, u_t) \leq 3(\hat{c}\mathcal{T})$, we have $T < \hat{c}\mathcal{T}$.

We prove Lemma 7 and 8 in supplementary material Section C. We also prove, in the same section, the main theorem using the coupling strategy of Jin et al. (2017a). but with the additional difficulty of taking into consideration the effect of the Riemannian geometry (Lemma 2) and the injectivity radius.

7 Examples

kPCA. We consider the kPCA problem, where we want to find the $k \leq n$ principal eigenvectors of a symmetric matrix $H \in \mathbb{R}^{n \times n}$, as an example (Tripurani et al., 2018). This corresponds to

$$\min_{X \in \mathbb{R}^{n \times k}} -\frac{1}{2} \text{tr}(X^T H X) \quad \text{subject to } X^T X = I,$$

which is an optimization problem on the Grassmann manifold defined by the constraint $X^T X = I$. If the eigenvalues of H are distinct, we denote by v_1, \dots, v_n the eigenvectors of H , corresponding to eigenvalues with decreasing order. Let $V^* = [v_1, \dots, v_k]$ be the matrix with columns composed of the top k eigenvectors of H , then the local minimizers of the objective function are $V^* G$ for all unitary matrices $G \in \mathbb{R}^{k \times k}$. Denote also by $V = [v_{i_1}, \dots, v_{i_k}]$ the matrix with columns composed of k distinct eigenvectors, then the first order stationary points of the objective function (with Riemannian gradient being 0) are $V G$ for all unitary matrices $G \in \mathbb{R}^{k \times k}$. In our numerical experiment, we choose H to be a diagonal matrix $H = \text{diag}(0, 1, 2, 3, 4)$ and let $k = 3$. The Euclidean basis (e_i) are an eigenbasis of H and the first order stationary points of the objective function are $[e_{i_1}, e_{i_2}, e_{i_3}] G$ with distinct basis and G being unitary. The local minimizers are $[e_3, e_4, e_5] G$. We start the iteration at $X_0 = [e_2, e_3, e_4]$ and see in Fig. 3 the algorithm converges to a local minimum.

Burer-Monteiro approach for certain low rank problems. Following Boumal et al. (2016b), we consider, for $A \in \mathbb{S}^{d \times d}$ and $r(r+1)/2 \leq d$, the problem

$$\min_{X \in \mathbb{S}^{d \times d}} \text{trace}(AX), \quad \text{s.t. } \text{diag}(X) = 1, X \succeq 0, \text{rank}(X) \leq r.$$

We factorize X by $Y Y^T$ with an overparametrized $Y \in \mathbb{R}^{d \times p}$ and $p(p+1)/2 \geq d$. Then any local minimum of

$$\min_{Y \in \mathbb{R}^{d \times p}} \text{trace}(A Y Y^T), \quad \text{s.t. } \text{diag}(Y Y^T) = 1,$$

is a global minimum where $Y Y^T = X^*$ (Boumal et al., 2016b). Let $f(Y) = \frac{1}{2} \text{trace}(A Y Y^T)$. In the experiment, we take $A \in \mathbb{R}^{100 \times 20}$ being a sparse matrix that only the upper left 5×5 block is random and other entries are 0. Let the initial point $Y_0 \in \mathbb{R}^{100 \times 20}$, such that $(Y_0)_{i,j} = 1$ for $5j - 4 \leq i \leq 5j$ and $(Y_0)_{i,j} = 0$ otherwise. Then Y_0 is a saddle point. We see in Fig. 3 the algorithm converges to the global optimum.

Summary We have shown that for the constrained optimization problem of minimizing $f(x)$ subject to a manifold constraint as long as the function and the manifold are appropriately smooth, a perturbed Riemannian gradient descent algorithm will escape saddle points with a rate of order $1/\epsilon^2$ in the accuracy ϵ , polylog in manifold dimension d , and depends polynomially on the curvature and smoothness parameters.

A natural extension of our result is to consider other variants of gradient descent, such as the heavy ball method, Nesterov's acceleration, and the stochastic setting. The question is whether these algorithms with appropriate modification (with manifold constraints) would have a fast convergence to second-order stationary point (not just first-order stationary as studied in recent literature), and whether it is possible to show the relationship between convergence rate and smoothness of manifold.

⁶“smallest eigenvector” means the eigenvector corresponding to the smallest eigenvalue.




Figure 3: (a) kPCA problem. We start from an approximate saddle point, and it converges to a local minimum (which is also global minimum). (b) Burer-Monteiro approach Plot $f(Y) = \frac{1}{2}\text{trace}(AYY^T)$ versus iterations. We start from the saddle point, and it converges to a local minimum (which is also global minimum).

References

- Absil, P.-A., Mahony, R., and Sepulchre, R. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2009.
- Agarwal, N., Boumal, N., Bullins, B., and Cartis, C. Adaptive regularization with cubics on manifolds with a first-order analysis. *arXiv preprint arXiv:1806.00065*, 2018.
- Ambrose, W. and Singer, I. M. A theorem on holonomy. *Transactions of the American Mathematical Society*, 75(3):428–443, 1953.
- Bécigneul, G. and Ganea, O.-E. Riemannian adaptive optimization methods. In *International Conference on Learning Representations*, 2019.
- Bonnabel, S. Stochastic gradient descent on riemannian manifolds. *IEEE Transactions on Automatic Control*, 58(9):2217–2229, 2013.
- Boumal, N. and Absil, P.-a. Rtrmc: A riemannian trust-region method for low-rank matrix completion. In *Advances in neural information processing systems*, pp. 406–414, 2011.
- Boumal, N., Mishra, B., Absil, P.-A., and Sepulchre, R. Manopt, a Matlab toolbox for optimization on manifolds. *JMLR*, 2014.
- Boumal, N., Absil, P.-A., and Cartis, C. Global rates of convergence for nonconvex optimization on manifolds. *IMA Journal of Numerical Analysis*, 2016a.
- Boumal, N., Voroninski, V., and Bandeira, A. The non-convex burer-monteiro approach works on smooth semidefinite programs. In *Advances in Neural Information Processing Systems*, pp. 2757–2765, 2016b.
- Boumal, N., Absil, P.-A., and Cartis, C. Global rates of convergence for nonconvex optimization on manifolds. *IMA Journal of Numerical Analysis*, pp. drx080, 2018. doi: 10.1093/imanum/drx080. URL <http://dx.doi.org/10.1093/imanum/drx080>.
- Carmon, Y. and Duchi, J. C. Gradient descent efficiently finds the cubic-regularized non-convex newton step. *arXiv preprint arXiv:1612.00547*, 2017.
- Cheeger, J. and Ebin, D. G. *Comparison Theorems in Riemannian Geometry*. AMS Chelsea Publishing, Providence, RI, 2008.
- Do Carmo, M. P. *Differential Geometry of Curves and Surfaces*. Courier Dover Publications, 2016.
- Du, S. S., Jin, C., Lee, J. D., Jordan, M. I., Singh, A., and Póczos, B. Gradient descent can take exponential time to escape saddle points. In *Advances in Neural Information Processing Systems*, pp. 1067–1077, 2017.
- Edelman, A., Arias, T. A., and Smith, S. T. The geometry of algorithms with orthogonality constraints. *SIAM journal on Matrix Analysis and Applications*, 20(2):303–353, 1998.

- Ge, R., Huang, F., Jin, C., and Yuan, Y. Escaping from saddle points – online stochastic gradient for tensor decomposition. In *Conference on Learning Theory*, pp. 797–842, 2015.
- Hu, J., Milzarek, A., Wen, Z., and Yuan, Y. Adaptive quadratically regularized Newton method for Riemannian optimization. *SIAM J. Matrix Anal. Appl.*, 39(3):1181–1207, 2018.
- Ishteva, M., Absil, P.-A., Van Huffel, S., and De Lathauwer, L. Best low multilinear rank approximation of higher-order tensors, based on the riemannian trust-region scheme. *SIAM Journal on Matrix Analysis and Applications*, 32(1):115–135, 2011.
- Jin, C., Ge, R., Netrapalli, P., Kakade, S., and Jordan, M. I. How to escape saddle points efficiently. In *ICML*, 2017a.
- Jin, C., Netrapalli, P., and Jordan, M. I. Accelerated gradient descent escapes saddle points faster than gradient descent. *arXiv preprint arXiv:1711.10456*, 2017b.
- Karcher, H. Riemannian center of mass and mollifier smoothing. *Communications on pure and applied mathematics*, 30(5):509–541, 1977.
- Kasai, H. and Mishra, B. Inexact trust-region algorithms on riemannian manifolds. In *Advances in Neural Information Processing Systems 31*, pp. 4254–4265. 2018.
- Khuzani, M. B. and Li, N. Stochastic primal-dual method on riemannian manifolds with bounded sectional curvature. *arXiv preprint arXiv:1703.08167*, 2017.
- Lee, J. D., Simchowitz, M., Jordan, M. I., and Recht, B. Gradient descent only converges to minimizers. *Conference on Learning Theory*, pp. 1246–1257, 2016.
- Lee, J. D., Panageas, I., Piliouras, G., Simchowitz, M., Jordan, M. I., and Recht, B. First-order methods almost always avoid saddle points. *arXiv preprint arXiv:1710.07406*, 2017.
- Lee, J. M. *Riemannian manifolds : an introduction to curvature*. Graduate texts in mathematics ; 176. Springer, New York, 1997. ISBN 9780387227269.
- Mangoubi, O., Smith, A., et al. Rapid mixing of geodesic walks on manifolds with positive curvature. *The Annals of Applied Probability*, 28(4):2501–2543, 2018.
- Miolane, N., Mathe, J., Donnat, C., Jorda, M., and Pennec, X. geomstats: a python package for riemannian geometry in machine learning. *arXiv preprint arXiv:1805.08308*, 2018.
- Mokhtari, A., Ozdaglar, A., and Jadbabaie, A. Escaping saddle points in constrained optimization. *arXiv preprint arXiv:1809.02162*, 2018.
- Nouiehed, M., Lee, J. D., and Razaviyayn, M. Convergence to second-order stationarity for constrained non-convex optimization. *arXiv preprint arXiv:1810.02024*, 2018.
- Pemantle, R. Nonconvergence to unstable points in urn models and stochastic approximations. *The Annals of Probability*, pp. 698–712, 1990.
- Rapcsák, T. Sectional curvatures in nonlinear optimization. *Journal of Global Optimization*, 40(1-3): 375–388, 2008.
- Sakai, T. *Riemannian Geometry*, volume 149 of *Translations of Mathematical Monographs*. American Mathematical Society, 1996.
- Sun, J., Qu, Q., and Wright, J. Complete dictionary recovery over the sphere ii: Recovery by riemannian trust-region method. *IEEE Transactions on Information Theory*, 63(2):885–914, 2017.
- Sun, Y. and Fazel, M. Escaping saddle points efficiently in equality-constrained optimization problems. In *Workshop on Modern Trends in Nonconvex Optimization for Machine Learning, International Conference on Machine Learning*, 2018.
- Tripuraneni, N., Flammarion, N., Bach, F., and Jordan, M. I. Averaging Stochastic Gradient Descent on Riemannian Manifolds. *arXiv preprint arXiv:1802.09128*, 2018.

- Wong, Y.-c. Sectional curvatures of Grassmann manifolds. *Proc. Nat. Acad. Sci. U.S.A.*, 60:75–79, 1968.
- Zhang, H. and Sra, S. First-order methods for geodesically convex optimization. *arXiv:1602.06053*, 2016. *Preprint*.
- Zhang, H., Reddi, S. J., and Sra, S. Riemannian svrg: fast stochastic optimization on riemannian manifolds. In *Advances in Neural Information Processing Systems*, pp. 4592–4600, 2016.
- Zhang, J. and Zhang, S. A cubic regularized newton’s method over riemannian manifolds. *arXiv preprint arXiv:1805.05565*, 2018.