

Geodesically convex M -estimation in metric spaces

Victor-Emmanuel Brunel

ENSAE

VICTOR.EMMANUEL.BRUNEL@ENSAE.FR

Editors: Gergely Neu and Lorenzo Rosasco

Abstract

We study the asymptotic properties of geodesically convex M -estimation on non-linear spaces. Namely, we prove that under very minimal assumptions besides geodesic convexity of the cost function, one can obtain consistency and asymptotic normality, which are fundamental properties in statistical inference. Our results extend the Euclidean theory of convex M -estimation; They also generalize limit theorems on non-linear spaces which, essentially, were only known for barycenters, allowing to consider robust alternatives that are defined through non-smooth M -estimation procedures.

Keywords: M -estimation, metric spaces, Riemannian manifolds, CAT spaces, geodesic convexity, barycenters, robust location estimation

1. Preliminaries

1.1. Introduction

The problem of M -estimation, or empirical risk minimization, is pervasive to statistics and machine learning. In many problems, one seeks to minimize a function that takes the form of an expectation, i.e., $\Phi(x) = \mathbb{E}[\phi(Z, x)]$, where Z is a random variable with unknown distribution, x is the target variable, which lives in some target space, and ϕ can be seen as a cost function. Classification, regression, location estimation, maximum likelihood estimation, are standard instances of such problems. In practice, the distribution of Z is unknown, so the function Φ cannot be evaluated. Hence, one rather minimizes a proxy of the form $\hat{\Phi}_n(x) = n^{-1} \sum_{i=1}^n \phi(Z_i, x)$, where Z_1, \dots, Z_n are available i.i.d. copies of Z . This method is called M -estimation, or empirical risk minimization. When the target space is Euclidean, a lot of results are available for guarantees of a minimizer \hat{x}_n of $\hat{\Phi}_n$, as an estimator for a minimizer x^* of Φ . Asymptotic results (such as consistency and asymptotic normality) are classic and well known, especially under technical smoothness assumptions on the cost function ϕ (Van Der Vaart and Wellner, 1996), most of which can be dropped when the cost function is convex in the target variable (Haberman, 1989; Niemiro, 1992).

In this work, we are interested in cases where the target space is not Euclidean, and does not even exhibit any linear structure. Of course, and at the very least, the space needs to be equipped with a metric in order to measure the accuracy of an estimation procedure. Thus, we consider a target space that is a metric space, which we denote by (M, d) .

Statistics and machine learning are more and more confronted with data that lie in non-linear spaces: In spatial statistics (e.g., directional data), computational tomography (e.g., data in quotient spaces such as in shape statistics, collected up to rigid transformations), economics (e.g., optimal transport, where data are measures), etc. Moreover, data and/or target parameters that are encoded as very high dimensional vectors may have a much smaller intrinsic dimension: For instance, they may lie on small dimensional submanifolds of the ambient Euclidean space. In that case, leveraging the

possibly non-linear geometry of the problems at hand can be a powerful tool in order to significantly reduce their dimensionality. This is one of the main motivations behind manifold learning (Genovese et al., 2012; Aamari and Levrard, 2019) or generative adversarial networks (Schreuder et al., 2021). Even though more and more algorithms are tailored to non-linear spaces (Lim and Pálfia, 2014; Ohta and Pálfia, 2015; Zhang and Sra, 2016, 2018; Antonakopoulos et al., 2020; Criscitiello and Boumal, 2022), still little is understood from a statistical prospective.

A specific instance of M -estimation on metric spaces, which has attracted much attention, is that of barycenters. Given n points $x_1, \dots, x_n \in M$, a barycenter is any minimizer $x \in M$ of the sum of squared distances $\sum_{i=1}^n d(x, x_i)^2$. One can easily check that if (M, d) is a Euclidean or Hilbert space, then the minimizer is unique and it is given by the linear average of x_1, \dots, x_n . Barycenters are a natural extension of linear averages to spaces with no linear structure. More generally, given a probability measure P on M , one can define a barycenter of P as any minimizer $x \in M$ of $\int_M d(x, z)^2 dP(z)$, provided that integral is finite for at least one (and hence, for all) value of $x \in M$. Barycenters were initially introduced in statistics by (Fréchet, 1948) in the 1940's, and later by (Karcher, 1977), and they were also known as Fréchet or Karcher means. They were popularized in the fields of shape statistics (Kendall et al., 2009), optimal transport (Agueh and Carlier, 2011; Cuturi and Doucet, 2014; Le Gouic and Loubes, 2017; Clatici et al., 2018; Kroshnin et al., 2019, 2021; Altschuler and Boix-Adsera, 2021, 2022). and matrix analysis (Bhatia and Holbrook, 2006; Bhatia, 2009; Bhatia et al., 2019). Asymptotic theory is fairly well understood for empirical barycenters in various setups, particularly laws of large numbers (Ziezold, 1977; Sturm, 2003) and central limit theorems in Riemannian manifolds (it is natural to impose a smooth structure on M in order to derive central limit theorems) (Bhattacharya and Patrangenaru, 2003, 2005; Kendall and Le, 2011; Huckemann, 2011; Bhattacharya and Lin, 2017; Eltzner and Huckemann, 2019; Eltzner et al., 2019). A few finite sample guarantees are available, even though the non-asymptotic statistical theory for barycenters is still quite limited (Sturm, 2003; Funano, 2010; Schötz, 2019; Ahidar-Coutrix et al., 2020; Le Gouic et al., 2022). Asymptotic theorems are also available for more general M -estimators on Riemannian manifolds, e.g., p -means, where $\mathcal{Z} = X$ and $\phi(z, x) = d(x, z)^p$, for some $p \geq 1$ (Schötz, 2019), including geometric medians as a particular case ($p = 1$), but only suboptimal rates are known. In particular, the standard $n^{-1/2}$ -rate is unknown in general, beyond barycenters. It should be noted that all central limit theorems for barycenters in Riemannian manifolds rely on the smoothness of the cost function $\phi = d^2$ and use standard techniques from smooth M -estimation (Van Der Vaart and Wellner, 1996) by performing Taylor expansions in local charts. Moreover, they assume that the data are almost surely in a ball with small radius, which essentially ensures the geodesic convexity of the squared distance function d^2 in its second variable, even though this synthetic property is not leveraged in this line of work. For instance, asymptotic normality cannot be obtained for geometric medians or other robust alternatives to barycenters, using these techniques.

1.2. Contributions

In this work, we prove strong consistency and asymptotic normality of geodesically convex M -estimators under very mild assumptions. We recover the central limit theorems proven for barycenters in Riemannian manifolds with small diameter, but we cover a much broader class of location estimators, including robust alternatives to barycenters. Just as in the Euclidean case, our results

convey that convexity is a powerful property, since it yields strong learning guarantees under very little assumptions.

All our results are asymptotic, but it is worth mentioning that proving learning guarantees in non-linear spaces is a very challenging, ongoing research topic, because it requires non-trivial tools from metric and differential geometry. In particular, very few non-asymptotic learning guarantees are available, even for simple estimators such as barycenters, and asymptotic theory is still an active research area in that framework (e.g., understanding non-standard asymptotic rates, a.k.a. sticky and smeary central limit theorems (Hotz et al., 2013; Eltzner and Huckemann, 2019)). Moreover, asymptotic guarantees such as asymptotic normality provide benchmarks that are also important for a finite sample theory.

2. Background and notation

2.1. Geodesic spaces

Let (M, d) be a complete and separable metric space that satisfies the following property: For all $x, y \in M$, there exists a continuous function $\gamma : [0, 1] \rightarrow M$ such that $\gamma(0) = x$, $\gamma(1) = y$ and with the property that $d(\gamma(s), \gamma(t)) = |s - t|d(x, y)$, for all $s, t \in [0, 1]$. Such a function γ is called a unit speed geodesic from x to y and it should be thought of as a (not necessarily unique) shortest path from x to y . We denote by $\Gamma_{x,y}$ the collection of all such functions. The space (M, d) is then called a geodesic space. Examples of geodesic spaces include Euclidean spaces, Euclidean spheres, finite and connected metric graphs, Wasserstein spaces, quotient spaces, etc. (Burago et al., 2001; Bridson and Haefliger, 2013). Geodesic spaces allow for a natural extension of the notion of convexity.

Definition 1 A function $f : M \rightarrow \mathbb{R}$ is called *geodesically convex* (convex, for short) if for all $x, y \in M, \gamma \in \Gamma_{x,y}$ and $t \in [0, 1]$, $f(\gamma(t)) \leq (1 - t)f(x) + tf(y)$.

In this work, we will further assume that (M, d) is a proper space, i.e., all bounded closed sets are compact. This assumption may seem limiting in practice but, at a high level, it essentially only discards infinite dimensional spaces. Indeed, Hopf-Rinow theorem (Bridson and Haefliger, 2013, Chapter 1, Proposition 3.7) asserts that so long as (M, d) is complete and locally compact, then it is proper. The version of Hopf-Rinow theorem for Riemannian manifolds asserts that any complete (finite dimensional) Riemannian manifold is a proper geodesic metric space (Do Carmo, 1992, Chapter 7, Theorem 2.8).

2.2. Riemannian manifolds

A Riemannian manifold (M, g) is a finite dimensional smooth manifold M equipped with a Riemannian metric g , i.e., a smooth family of scalar products on tangent spaces. We refer the reader to (Do Carmo, 1992) and (Lee, 2012, 2018) for a clear introduction to differential geometry and Riemannian manifolds. The distance inherited on M from the Riemannian metric g will be denoted by d . In this work, for simplicity, we only consider Riemannian manifolds without boundary, even though our results can be easily extended to manifolds with boundary.

The tangent bundle of M is denoted by TM , i.e., $TM = \bigcup_{x \in M} T_x M$ where, for all $x \in M$, $T_x M$ is the tangent space to M at x . For all $x \in M$, we let $\text{Exp}_x : T_x M \rightarrow M$ and $\text{Log}_x : M \rightarrow T_x M$

be the exponential and the logarithmic maps at x , respectively. The latter may not be defined on the whole tangent space, but it is always defined at least on a neighborhood of the origin. To provide intuition, at any given point $x \in M$, the tangent space $T_x M$ can be seen as a linear space attached to x , where each vector indicates a direction where to move on M starting from x : For a given vector $u \in T_x M$, $\text{Exp}_x(u)$ is the Riemannian analog of “ $x + u$ ” in the Euclidean space: From point x , add a vector u , whose direction indicates in which direction to move, and whose norm (given by the Riemannian metric g) indicates at which distance, in order to attain the point on M that is denoted by $\text{Exp}_x(u)$. On the opposite, given $x, y \in M$, $\text{Log}_x(y)$ (when it is well-defined) is the Riemannian analog of “ $y - x$ ”, that is, from x , which vector $u \in T_x M$ (if any) led to move to y , i.e., satisfies $\text{Exp}_x(u) = y$.

For all $x \in M$, we denote by $\langle \cdot, \cdot \rangle_x$ the scalar product on $T_x M$ inherited from g and by $\| \cdot \|_x$ the corresponding norm (i.e., for all $u, v \in T_x M$, $\langle u, v \rangle_x = g_x(u, v)$ and $\|u\|_x = \sqrt{g_x(u, u)}$).

Let $x, y \in M$ and $\gamma \in \Gamma_{x,y}$ be a geodesic (in the sense of Section 2.1) from x to y . If x and y are close enough, γ is unique (Lee, 2018, Proposition 6.11) and in that case, for all $t \in [0, 1]$, $\gamma(t) = \text{Exp}_x(t\dot{\gamma}(0))$, where $\dot{\gamma}(0) \in T_x M$ is the derivative of γ at $t = 0$, given by $\dot{\gamma}(0) = \text{Log}_x(y)$. We denote the parallel transport map from x to y along γ as $\pi_{x,y} : T_x M \rightarrow T_y M$, without specifying its dependence on γ (which is non-ambiguous if y is close enough to x). The map $\pi_{x,y}$ is a linear isometry, i.e., $\langle u, v \rangle_x = \langle \pi_{x,y}(u), \pi_{x,y}(v) \rangle_y$, for all $u, v \in T_x M$. Moreover, it holds that $\pi_{x,y}(\dot{\gamma}(0)) = -\dot{\gamma}(1)$. Intuitively, parallel transport is an important tool that allows to compare vectors from different tangent spaces, along a given path on the manifold.

For further details and properties on Riemannian manifolds used in the mathematical development of this work, we refer the reader to Appendix B.

2.3. M -estimation

Let \mathcal{Z} be some abstract space, equipped with a σ -algebra \mathcal{F} and a probability measure P . Let $\phi : \mathcal{Z} \times M \rightarrow \mathbb{R}$ satisfy the following:

- (i) For all $x \in M$, the map $\phi(\cdot, x)$ is measurable and integrable with respect to P ;
- (ii) For P -almost all $z \in \mathcal{Z}$, $\phi(z, \cdot)$ is convex.

Finally, let Z, Z_1, Z_2, \dots be i.i.d. random variables in \mathcal{Z} with distribution P and set the functions

$$\Phi(x) = \mathbb{E}[\phi(Z, x)] \quad \text{and} \quad \hat{\Phi}_n(x) = n^{-1} \sum_{i=1}^n \phi(Z_i, x),$$

for all $x \in M$ and for all positive integers n . Finally, we denote by M^* the set of minimizers of Φ . Note that by convexity of Φ , the minimizing set M^* is convex, i.e., for all $x, y \in M^*$ and for all $\gamma \in \Gamma_{x,y}$, $\gamma([0, 1]) \subseteq M^*$.

Important examples include the case where $\mathcal{Z} = M$ and ϕ is a function of the metric, which yields location estimation. Specifically, assume that $\mathcal{Z} = M$ and that $\phi(z, x) = \ell(d(z, x))$, for all $z, x \in M$, where $\ell : [0, \infty) \rightarrow [0, \infty)$ is a non-decreasing, convex function. Then, a natural framework to ensure convexity of the functions $\phi(z, \cdot)$ is that of spaces with curvature upper bounds, a.k.a. CAT spaces. Here, we only give an intuitive definition on CAT spaces. We refer the reader to Appendix D for a more detailed account on CAT spaces.

Fix some real number κ . We say that (M, d) is $\text{CAT}(\kappa)$ (or that it has global curvature upper bound κ) if all small enough triangles in M are thinner than what they would be in a space of constant curvature κ , i.e., a sphere if $\kappa > 0$, a Euclidean space if $\kappa = 0$ and a hyperbolic space if $\kappa < 0$. This setup is very natural to grasp the convexity of the distance function to a point. Fix $x_0 \in M$. The function $d(\cdot, x_0)$ is convex if and only if for all $y, z \in M$ and all $\gamma \in \Gamma_{y,z}$, $d(\gamma(t), x_0) \leq (1-t)d(y, x_0) + td(z, x_0)$, which controls the distance from x_0 to any point in the opposite edge to x_0 in a triangle with vertices x_0, y, z . In other words, it indicates how thin a triangle with vertices x_0, y, z must be.

Hence, if $\mathcal{Z} = M$ is $\text{CAT}(\kappa)$ and has diameter at most D_κ (see Appendix D for its definition), (Ohta, 2007, Proposition 3.1) guarantees that for all $z \in M$, $d(z, \cdot)$ is convex on M and hence, $\phi(z, \cdot) = \ell(\phi(z, \cdot))$ is also convex as soon as $\ell : [0, \infty) \rightarrow [0, \infty)$ is non-decreasing and convex itself. In that setup, important examples of functions ℓ include:

- (i) $\ell(u) = u^2$: Then, a minimizer of Φ is a barycenter of P and a minimizer of $\hat{\Phi}_n$ is an empirical barycenter of Z_1, \dots, Z_n
- (ii) $\ell(u) = u$: Then, a minimizer of Φ is a geometric median of P
- (iii) More generally, if $\ell(u) = u^p, p \geq 1$, a minimizer of Φ is called a p -mean of P
- (iv) $\ell(u) = u^2$ if $0 \leq u \leq c$, $\ell(u) = c(2u - c)$ if $u > c$, where $c > 0$ is a fixed parameter. This function was introduced by Huber (Huber, 1992) in order to produce low-bias robust estimators of the mean: It provides an interpolation between barycenter and geometric median estimation.

Important examples of CAT spaces include Euclidean spaces, spheres, simply connected Riemannian manifolds with sectional curvature bounded from above, metric trees, etc. In Appendix D, we provide a more detailed list of examples.

3. Consistency

This section contains our first main results on the consistency of geodesically convex M -estimators. Here, we work under two scenarios. First, we assume that $\phi(z, \cdot)$ is convex and Lipschitz with some constant that does not depend on $z \in \mathcal{Z}$. This is somewhat restrictive, even though it covers most of the cases mentioned above for location estimation. In the second scenario, we only assume convexity, but we consider the case where M is a Riemannian manifold. We then discuss possible extensions of our findings.

3.1. Lipschitz case

Here, we assume that there exists $L > 0$ such that $\phi(z, \cdot)$ is L -Lipschitz on M , for all $z \in \mathcal{Z}$. For instance, this is satisfied if M has bounded diameter and $\phi = \ell \circ d$, for some locally Lipschitz function $\ell : [0, \infty)$, such as the functions mentioned above for location estimation.

Recall that we denote by M^* the set of minimizers of Φ . The following theorem need not require that Φ has a unique minimizer: It asserts that any minimizer of $\hat{\Phi}_n$ will eventually become arbitrarily close to M^* with probability 1.

Theorem 2 *Let M be a proper geodesic metric space. Let $\phi(z, \cdot)$ be geodesically convex and L -Lipschitz for P -almost all $z \in \mathcal{Z}$, where $L > 0$ is some positive number. Assume that M^* is non-empty and bounded. For $n \geq 1$, let \hat{x}_n be a minimizer of $\hat{\Phi}_n$. Then, it holds that $d(\hat{x}_n, M^*) \xrightarrow[n \rightarrow \infty]{} 0$ almost surely.*

A simple adaptation of the proof shows that the same conclusion applies if \hat{x}_n is not exactly a minimizer of $\hat{\Phi}_n$ but rather satisfies $\hat{\Phi}_n(\hat{x}_n) \leq \min_{x \in M} \hat{\Phi}_n(x) + \varepsilon_n$, where ε_n is any non-negative error term that goes to zero almost surely, as $n \rightarrow \infty$.

The proof of Theorem 2 is based on the following extension of (Rockafellar, 1970, Theorem 10.8). The proof is a straightforward adaptation of that of (Rockafellar, 1970, Theorem 10.8).

Lemma 3 *Let $L > 0$ and $(f_n)_{n \geq 1}$ be a sequence of L -Lipschitz geodesically convex functions on M . Assume that there is a function $f : M \rightarrow \mathbb{R}$ and a dense subset M_0 of M such that $f_n(x) \xrightarrow[n \rightarrow \infty]{} f(x)$, for all $x \in M_0$. Then, f is geodesically convex, L -Lipschitz, and f_n converges to f uniformly on any compact subset of M .*

This lemma is the reason why, in this section, we require $\phi(z, \cdot)$ to be Lipschitz, with a constant that does not depend on $z \in \mathcal{Z}$. If M was Euclidean, the Lipschitz assumption of Lemma 3 would not be necessary because on any compact subset, it would be a consequence of the convexity and pointwise convergence of the functions $f_n, n \geq 1$, see (Rockafellar, 1970, Theorem 10.6). However, to the best of our knowledge, this may not hold in the more general case that we treat here (and may be subject to a further research question).

As a corollary to Lemma 3, we now state the following results, which is essential in our proof of Theorem 2. The proof is deferred to Appendix C.

Lemma 4 *Let $L > 0$ and $(F_n)_{n \geq 1}$ be a sequence of real valued, L -Lipschitz, geodesically convex random functions on M . Let F a (possibly random) function on M and assume that for all $x \in M$, $F_n(x)$ converges almost surely (resp. in probability) to $F(x)$. Then, the convergence holds uniformly on any compact subset K of M , i.e., $\sup_{x \in K} |F_n(x) - F(x)|$ converges to zero almost surely (resp. in probability).*

PROOF OF THEOREM 2. Let $\Phi^* = \min_{x \in M} \Phi(x)$ be the smallest value of Φ on M and fix some arbitrary $x^* \in M^*$. Fix $\varepsilon > 0$ and let $K_\varepsilon = \{x \in M : d(x, M^*) = \varepsilon\}$. Since M^* is bounded, K_ε is bounded. Moreover, since the distance is continuous, K_ε is a closed set. Hence, it is compact, since we have assumed (M, d) to be proper. Since Φ is Lipschitz, it is continuous so it holds that $\eta := \min_{x \in K_\varepsilon} \Phi(x) - \Phi^* > 0$. Moreover, by the law of large numbers, $\Phi_n(x) \xrightarrow[n \rightarrow \infty]{} \Phi(x)$ almost surely, for all $x \in M$. Therefore, by Lemma 4, Φ_n converges uniformly to Φ on K_ε almost surely. So, with probability 1, $\inf_{x \in K_\varepsilon} \hat{\Phi}_n(x) > \Phi^* + 2\eta/3$ for all large enough n . Moreover, also with probability 1, $\hat{\Phi}_n(x^*) < \Phi^* + \eta/3$ for all large enough n .

We have established that with probability 1, it holds simultaneously, for all large enough n , that $\inf_{x \in K_\varepsilon} \hat{\Phi}_n(x) > \hat{\Phi}_n(x^*)$. Let us show that this, together with the convexity of $\hat{\Phi}_n$, implies that any minimizer of $\hat{\Phi}_n$ must be at a distance at most ε of M^* . This will yield the desired result.

Assume, for the sake of contradiction, that $\hat{\Phi}_n$ has a minimizer \tilde{x}_n that satisfies $d(\tilde{x}_n, M^*) > \varepsilon$. Consider a geodesic $\gamma \in \Gamma_{x^*, \tilde{x}_n}$. Then, by continuity of the function $d(\cdot, M^*)$, there must be some $t \in [0, 1]$ such that $\gamma(t) \in K_\varepsilon$. The convexity of $\hat{\Phi}_n$ yields the convexity of $\hat{\Phi}_n \circ \gamma$, which

is minimum at $t = 1$. Hence, $\hat{\Phi}_n \circ \gamma$ must be non-increasing, implying that $\inf_{x \in K_\varepsilon} \hat{\Phi}_n(x) \leq \hat{\Phi}_n(\gamma(t)) \leq \hat{\Phi}_n(x^*)$, which yields a contradiction. \square

3.2. Riemannian framework

Now, we prove that, at least in a Riemannian manifold, the Lipschitz condition on ϕ can be dropped for the consistency of \hat{x}_n . In this section, we assume that (M, g) is a complete Riemannian manifold and we have the following theorem.

Theorem 5 *Let (M, g) be a Riemannian manifold. Let $\phi(z, \cdot)$ be geodesically convex for P -almost all $z \in \mathcal{Z}$. Assume that M^* is non-empty and bounded. For $n \geq 1$, let \hat{x}_n be a minimizer of $\hat{\Phi}_n$. Then, it holds that $d(\hat{x}_n, M^*) \xrightarrow{n \rightarrow \infty} 0$ almost surely.*

Again, this theorem applies if \hat{x}_n satisfies $\hat{\Phi}_n(\hat{x}_n) \leq \min_{x \in M} \hat{\Phi}_n(x) + \varepsilon_n$, where ε_n is any non-negative error term that goes to zero almost surely, as $n \rightarrow \infty$.

Once one has Lemma 7 below, which builds upon Lemma 6, the proof of Theorem 5 is very similar to that of Theorem 2, hence, we omit it.

Lemma 6 (*Greene and Wu, 1973*) *Any geodesically convex function on a Riemannian manifold is continuous and locally Lipschitz.*

By adapting the proof of (Rockafellar, 1970, Theorem 10.8), this lemma yields the following result.

Lemma 7 *Let $(f_n)_{n \geq 1}$ be a sequence of geodesically convex functions on M . Assume that $f_n(x) \xrightarrow{n \rightarrow \infty} f(x)$, for all x in a dense subset M_0 of M , where $f : M \rightarrow \mathbb{R}$ is some given geodesically convex function. Then, f_n is equi-Lipschitz and converges to f uniformly on any compact subset of M .*

Remark 8 *Lemma 6 is key in the proof of consistency if one does not assume that $\phi(z, \cdot)$ is L -Lipschitz, for all $z \in \mathcal{Z}$, with some $L > 0$ that does not depend on z . As we pointed out in the previous section, we do not know whether this lemma still holds true in more general proper geodesic spaces and we leave this as an open question.*

4. Asymptotic normality

Now, we turn to asymptotic normality of \hat{x}_n . Riemannian manifolds offer a natural and reasonable framework because, on top of their metric structure, they enjoy smoothness which allows to compute first and second order expansions and Gaussian distributions can be defined on their tangent spaces. Hence, in this section, we assume that (M, g) be a complete Riemannian manifold.

Theorem 9 *Let $\phi(z, \cdot)$ be geodesically convex, for P -almost all $z \in \mathcal{Z}$. Assume that the distribution P and the function ϕ satisfy the following assumptions:*

- Φ has a unique minimizer $x^* \in M$
- Φ is twice differentiable at x^* and $H\Phi(x^*)$ is positive definite

- There exists $\eta > 0$ such that $\mathbb{E}[\|g(Z, x)\|_{x^*}^2] < \infty$, for all $x \in M$ with $d(x^*, x) \leq \eta$, where $g(Z, x)$ is a measurable subgradient of $\phi(Z, \cdot)$ at the point x .

Then, \hat{x}_n is asymptotically normal, that is,

$$\sqrt{n}\text{Log}_{x^*}(\hat{x}_n) \xrightarrow[n \rightarrow \infty]{(d)} \mathcal{N}_{T_{x^*}M}(0, V(x^*))$$

where $V(x^*) = S(x^*)^{-1}B(x^*)S(x^*)^{-1}$, $S(x^*) = H\Phi(x^*)$, $B(x^*) = \mathbb{E}[g(Z, x^*)g(Z, x^*)^\top]$ and we denote by $\mathcal{N}_{T_{x^*}M}(0, V(x^*))$ the normal distribution on the Euclidean space $T_{x^*}M$ with mean 0 and covariance operator $V(x^*)$.

Here, $H\Phi(x^*)$ is the Hessian of Φ at the point x^* (see Appendix B). Recall that $\text{Log}_{x^*}(\hat{x}_n)$ is the Riemannian analog of “ $\hat{x}_n - x^*$ ”, so Theorem 9 is a natural formulation of asymptotic normality of \hat{x}_n in the Riemannian context.

A close look to the end of the proof will convince the reader that one can assume, without modifying the conclusion of this theorem, that \hat{x}_n need not be a minimizer of $\hat{\Phi}_n$, but should satisfy that $\hat{\Phi}_n(\hat{x}_n) \leq \min_{x \in M} \hat{\Phi}_n(x) + o_P(n^{-1/2})$, where $o_P(n^{-1/2})$ is a possibly random error term that goes to zero in probability at a faster rate than $n^{-1/2}$.

Remark 10 *The statement of Theorem 9 implicitly assumes the existence of measurable subgradients of ϕ , that is, a mapping $g : \mathcal{Z} \times M \rightarrow TM$ such that for P -almost all $z \in \mathcal{Z}$ and for all $x \in M$, $g(z, x) \in T_xM$ is a subgradient of $\phi(z, \cdot)$ at the point $x \in M$, and for all $x \in M$, $g(\cdot, x) : \mathcal{Z} \rightarrow T_xM$ is measurable. This is actually guaranteed by Lemma 14 on measurable selections, see Appendix A. Note also that the last assumption is automatically guaranteed if $\phi(z, \cdot)$ is locally $L(z)$ -Lipschitz at x^* , for P -almost all $z \in \mathcal{Z}$, for some $L \in L^2(P)$.*

The proof of Theorem 9 is strongly inspired by (Haberman, 1989) and (Niemiro, 1992) but requires subtle tools to account for the non-Euclidean geometry of the problem. The idea of the proof is to show that

$$\hat{\Phi}_n(\text{Exp}_{x^*}(u/\sqrt{n})) \approx \hat{\Phi}_n(x^*) + \frac{1}{\sqrt{n}} \langle u, n^{-1} \sum_{i=1}^n g(Z_i, x^*) \rangle_{x^*} + n^{-1} \langle u, H\Phi(x^*)u \rangle_{x^*}$$

as n grows large, uniformly in $u \in T_{x^*}M$. Hence, the minimizer of the left-hand side which, by definition, is $\text{Log}_{x^*}\hat{x}_n$, must be close to the minimizer of the right-hand side, which has a closed form, and which is asymptotically normal, thanks to the central limit theorem. The main difficulty is to handle $\hat{\Phi}_n(\text{Exp}_{x^*}(u/\sqrt{n}))$, which is not a convex function of $u \in T_{x^*}M$, even though $\hat{\Phi}_n$ is convex on M . Lemma 11 below is our most technical result, and is the key argument to adapt the techniques of (Haberman, 1989; Niemiro, 1992) to the Riemannian setup.

Lemma 11

Let (M, g) be a Riemannian manifold and $x_0 \in M$. There exists $r > 0$ such that the following holds. Let $L > 0$. There exists $\alpha > 0$ such that for all geodesically convex functions $f : M \rightarrow \mathbb{R}$ that are L -Lipschitz on $B(x_0, r)$, the map

$$f \circ \text{Exp}_{x_0} + \frac{\alpha}{2} \|\cdot\|_{x_0}^2 : T_{x_0}M \rightarrow \mathbb{R}$$

is convex on $\{u \in T_{x_0}M : \|u\|_{x_0} \leq r/2\}$, where $\|\cdot\|_{x_0}$ is the norm induced by g on the tangent space $T_{x_0}M$.

In the proof of this lemma, we first consider the case where f is smooth and then proceed by approximation of convex functions by smooth functions, leveraging (Greene and Wu, 1973, Theorem 2).

PROOF OF THEOREM 9. Fix $u \in T_{x^*}M$. For all $i = 1, \dots, n$, let

$$W_{i,n} = \phi(Z_i, \text{Exp}_{x^*}(u/\sqrt{n})) - \phi(Z_i, x^*) - \frac{1}{\sqrt{n}} \langle u, g(Z_i, x^*) \rangle_{x^*}.$$

The definition of subgradients yields that $W_{i,n} \geq 0$. Let π_n be the parallel transport from x^* to $x_n := \text{Exp}_{x^*}(u/\sqrt{n})$ through the geodesic $\gamma(t) = \text{Exp}_{x^*}(tu/\sqrt{n}), t \in [0, 1]$. Then, $x^* = \text{Exp}_{x_n}(-\pi_n^{-1}(u)/\sqrt{n})$. Thus, one has the following:

$$\begin{aligned} W_{i,n} &\leq \langle \pi_n(u)/\sqrt{n}, g(Z_i, x_n) \rangle_{x_n} - \langle u/\sqrt{n}, g(Z_i, x^*) \rangle_{x^*} \\ &= \langle u/\sqrt{n}, \pi_n^{-1}(g(Z_i, \text{Exp}_{x^*}(u/\sqrt{n}))) - g(Z_i, x^*) \rangle_{x^*}, \end{aligned} \quad (1)$$

where we used the fact that π_n is an isometry in the last equality.

Therefore, $\mathbb{E}[W_{i,n}^2] \leq n^{-1}\mathbb{E}[Y_n^2]$, where

$$Y_n = \langle u, \pi_n^{-1}(g(Z_1, \text{Exp}_{x^*}(u/\sqrt{n}))) - g(Z_1, x^*) \rangle_{x^*}.$$

Note that $\mathbb{E}[Y_n^2]$ is finite if n is large enough, thanks to the last assumption of the theorem.

The inequalities above yield that $Y_n \geq 0$ for all $n \geq 1$ and by Lemma 16, the sequence $(Y_n)_{n \geq 1}$ is non-increasing. Therefore, Y_n converges almost surely to a non-negative random variable Y . Let us now prove that $Y = 0$ almost surely. In order to do so, let us prove that $\mathbb{E}[Y_n] \xrightarrow{n \rightarrow \infty} 0$. Since, by the monotone convergence theorem, $\mathbb{E}[Y_n]$ must converge to $\mathbb{E}[Y]$, we will obtain that $\mathbb{E}[Y] = 0$. This, combined with the fact that $Y \geq 0$ almost surely, will yield the desired result.

Fix $x \in M$ with $d(x, x^*) \leq \eta$. Then, for all $v \in T_x M$, $\phi(Z_1, \text{Exp}_x(v)) \geq \phi(Z_1, x) + \langle v, g(Z_1, x) \rangle_x$. Taking the expectation readily yields that $\mathbb{E}[g(Z_1, x)] \in \partial\Phi(x)$. Hence, for all integers n that are large enough so $\|u/\sqrt{n}\|_{x^*} \leq \eta$, it holds that $g_n := \mathbb{E}[g(Z_1, \text{Exp}_{x^*}(u/\sqrt{n}))] \in \partial\Phi(\text{Exp}_{x^*}(u/\sqrt{n}))$. Therefore,

$$\mathbb{E}[Y_n] = \langle u, \pi_n^{-1}(g_n) - \nabla\Phi(x^*) \rangle_{x^*} \xrightarrow{n \rightarrow \infty} 0$$

since Φ is twice differentiable at x^* .

Finally, we obtain that $Y = 0$ almost surely. Now, $(Y_n)_{n \geq 0}$ is a non-increasing, non-negative sequence, so the monotone convergence theorem yields that $\mathbb{E}[Y_n^2]$ goes to zero as $n \rightarrow \infty$, so that

$$\text{var} \left(\sum_{i=1}^n W_{i,n} \right) = \sum_{i=1}^n \text{var}(W_{i,n}) \leq \sum_{i=1}^n \mathbb{E}[W_{i,n}^2] \leq \mathbb{E}[Y_n^2] \xrightarrow{n \rightarrow \infty} 0.$$

Therefore, by Chebychev's inequality, $\sum_{i=1}^n (W_{i,n} - \mathbb{E}[W_{i,n}])$ converges in probability to zero, as $n \rightarrow \infty$, that is,

$$\begin{aligned} n \left(\hat{\Phi}_n(\text{Exp}_{x^*}(u/\sqrt{n})) - \hat{\Phi}_n(x^*) \right) - \frac{1}{\sqrt{n}} \langle u, \sum_{i=1}^n g(Z_i, x^*) \rangle_{x^*} \\ - n \left(\Phi(\text{Exp}_{x^*}(u/\sqrt{n})) - \Phi(x^*) - \langle u, \nabla\Phi(x^*) \rangle_{x^*} \right) \xrightarrow{n \rightarrow \infty} 0 \end{aligned}$$

in probability. Since Φ is twice differentiable at x^* , we obtain

$$n \left(\hat{\Phi}_n(\text{Exp}_{x^*}(u/\sqrt{n})) - \hat{\Phi}_n(x^*) \right) - \frac{1}{\sqrt{n}} \langle u, \sum_{i=1}^n g(Z_i, x^*) \rangle_{x^*} - \langle u, \mathbf{H}\Phi(x^*)u \rangle_{x^*} \xrightarrow[n \rightarrow \infty]{} 0 \quad (2)$$

in probability. For $n \geq 1$ and $u \in T_{x^*}M$, denote by

$$F_n(u) = n \left(\hat{\Phi}_n(\text{Exp}_{x^*}(u/\sqrt{n})) - \hat{\Phi}_n(x^*) \right) - \frac{1}{\sqrt{n}} \langle u, \sum_{i=1}^n g(Z_i, x^*) \rangle_{x^*}$$

and by

$$F(u) = \langle u, \mathbf{H}\Phi(x^*)u \rangle_{x^*}.$$

We have shown that $F_n(u)$ converges in probability to $F(u)$ as $n \rightarrow \infty$, for any fixed $u \in T_{x^*}M$. Now, we would like to make use of Lemma 4. However, the functions F_n are not guaranteed to be convex. We adjust these functions using lemma 11.

Let $R > 0$ be any fixed positive number. Let $K = B_{x^*}(0, R)$ be the ball in $T_{x^*}M$ centered at the origin, with radius R and let $\tilde{K} = B(x^*, R)$. Then, K and \tilde{K} are both compact. By the law of large numbers, $\hat{\Phi}_n(x)$ converges almost surely to $\Phi(x)$, for all $x \in \tilde{K}$. Therefore, by Lemma 4, $\sup_{x \in \tilde{K}} |\hat{\Phi}_n(x) - \Phi(x)| \xrightarrow[n \rightarrow \infty]{} 0$ almost surely. Hence, Lemma 7 yields the existence of a (possibly random) $L > 0$ such that the following holds with probability 1:

$$|\hat{\Phi}_n(x) - \hat{\Phi}_n(y)| \leq Ld(x, y), \forall n \geq 1, \forall x, y \in \tilde{K}.$$

Now, let $r > 0$ be as in Lemma 11. Then, there exists a (possibly random) $\alpha > 0$ such that with probability 1, for all $n \geq 1$, $\hat{\Phi}_n \circ \text{Exp}_{x^*} + \frac{\alpha}{2} \|\cdot\|_{x^*}^2$ is convex on K .

Denote by $\tilde{F}_n(u) = F_n(u) + \frac{\alpha}{2} \|u\|_{x^*}^2$ and $\tilde{F}(u) = F(u) + \frac{\alpha}{2} \|u\|_{x^*}^2$, for all $u \in K$ and $n \geq 1$. For all large enough $n \geq 1$ (such that $R/\sqrt{n} \leq r$), each F_n is convex on K with probability 1, and we trivially have that $\tilde{F}_n(u) \xrightarrow[n \rightarrow \infty]{} \tilde{F}(u)$ in probability, for all $u \in K$. Therefore, Lemma 4 yields that $\sup_{u \in K} |F_n(u) - F(u)| = \sup_{u \in K} |\tilde{F}_n(u) - \tilde{F}(u)| \xrightarrow[n \rightarrow \infty]{} 0$ in probability.

Finally, we have shown that for any $R > 0$,

$$\sup_{\|u\|_{x^*} \leq R} \left| n \left(\hat{\Phi}_n(\text{Exp}_{x^*}(u/\sqrt{n})) - \hat{\Phi}_n(x^*) \right) - \frac{1}{\sqrt{n}} \langle u, \sum_{i=1}^n g(Z_i, x^*) \rangle_{x^*} - \langle u, \mathbf{H}\Phi(x^*)u \rangle_{x^*} \right| \xrightarrow[n \rightarrow \infty]{} 0 \quad (3)$$

in probability.

By definition of \hat{x}_n , $\hat{\Phi}_n(\text{Exp}_{x^*}(u/\sqrt{n}))$ is minimized for $u = \sqrt{n}\text{Log}_{x^*}(\hat{x}_n)$. Moreover, one easily checks that $\frac{1}{\sqrt{n}} \langle u, \sum_{i=1}^n g(Z_i, x^*) \rangle_{x^*} - \langle u, \mathbf{H}\Phi(x^*)u \rangle_{x^*}$ is minimized for

$$u = -\mathbf{H}\Phi(x^*)^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n g(Z_i, x^*).$$

Note that $g(Z_1, x^*), \dots, g(Z_n, x^*)$ are i.i.d. random vectors in $T_{x^*}M$. They are centered, since $\mathbb{E}[g(Z_1, x^*)] = \nabla\Phi(x^*) = 0$ by the first order condition, and they have a second moment, thanks to the last assumption of the theorem. Therefore, the multivariate central limit theorem yields that

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n g(Z_i, x^*) \xrightarrow[n \rightarrow \infty]{} \mathcal{N}_{T_{x^*}M}(0, B(x^*)) \quad (4)$$

in distribution. In particular, the sequence $\frac{1}{\sqrt{n}} \sum_{i=1}^n g(Z_i, x^*)$ is bounded in probability, so one can choose R in (3) so that $\frac{1}{\sqrt{n}} \sum_{i=1}^n g(Z_i, x^*) \in B_{x_0}(0, R)$ for all $n \geq 1$, with arbitrarily large probability.

From (3), it is easy to obtain that $\sqrt{n} \text{Log}_{x^*}(\hat{x}_n) - \text{H}\Phi(x^*)^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n g(Z_i, x^*)$ converges in probability to zero, as $n \rightarrow \infty$. Therefore, by Slutsky's theorem, $\sqrt{n} \text{Log}_{x^*}(\hat{x}_n)$ has the same limit in distribution as $\text{H}\Phi(x^*)^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n g(Z_i, x^*)$, which concludes the proof thanks to (4). \square

PROOF OF LEMMA 11. Fix $r > 0$ (to be chosen) and denote by $K = B(x_0, r)$, which is compact by Hopf-Rinow theorem (Do Carmo, 1992, Chapter 7, Theorem 2.8). Let $f : M \rightarrow \mathbb{R}$ be convex and L -Lipschitz ($L > 0$) on K and set $F = f \circ \text{Exp}_{x_0} : T_{x_0}M \rightarrow \mathbb{R}$.

We split the proof of this lemma into steps.

Step 1: Case where f is smooth. First, assume that f is smooth. Then, F is a smooth function. In the sequel, we denote by $\text{H}f$ the Hessian of f and by $\bar{\text{H}}F$ the Hessian of F (i.e., we use H for Hessian of functions defined on M and $\bar{\text{H}}$ for Hessian of functions defined on the tangent space $T_{x_0}M$). Fix $u, v \in T_{x_0}M$ and let $\phi(t) = F(u + tv)$, for $s \in \mathbb{R}$. Denote by $\eta(s, t) = \text{Exp}_{x_0}(t(u + sv))$, for all $s, t \in \mathbb{R}$. The Hessian of F is given by $\bar{\text{H}}F(u)(v, v) = \phi''(0)$. For all $s \in \mathbb{R}$,

$$\phi'(s) = \langle \nabla f(\eta(s, 1)), \frac{\partial \eta}{\partial s}(s, 1) \rangle_{\eta(s, 1)}$$

and

$$\phi''(s) = \langle \text{H}f(\eta(s, 1)) \frac{\partial \eta}{\partial s}(s, 1), \frac{\partial \eta}{\partial s}(s, 1) \rangle_{\eta(s, 1)} + \langle \nabla f(\eta(s, 1)), D_s \frac{\partial \eta}{\partial s}(s, 1) \rangle_{\eta(s, 1)} \quad (5)$$

where D_s is the covariant derivative along the path $\eta(\cdot, 1)$. Note that the first term on the right-hand side of (5) is non-negative, due to the convexity of f (see (Udriste, 2013, Theorem 6.2)). Therefore,

$$\begin{aligned} \phi''(0) &\geq \langle \nabla f(\eta(0, 1)), D_s \frac{\partial \eta}{\partial s}(0, 1) \rangle_{\eta(0, 1)} \geq -\|\nabla f(\eta(0, 1))\| \|D_s \frac{\partial \eta}{\partial s}(0, 1)\| \\ &\geq -L \|D_s \frac{\partial \eta}{\partial s}(0, 1)\| \end{aligned} \quad (6)$$

where the second inequality follows from Cauchy-Schwarz inequality and the last one follows from the assumption that f is L -Lipschitz on K .

For all $w, w' \in T_{x_0}M$, let $\Gamma_{w, w'}(s, t) = \text{Exp}_{x_0}(t(w + sw'))$, $s, t \in [0, 1]$. In particular, $\eta = \Gamma_{u, v}$. Then, $D_s \frac{\partial \Gamma_{u, v}}{\partial s}(0, 1) = \|v\|^2 D_s \frac{\partial \Gamma_{u, \tilde{v}}}{\partial s}(0, 1)$, where $\tilde{v} = \frac{v}{\|v\|_{x_0}}$. Since the map $w, w' \mapsto D_s \frac{\partial \Gamma_{w, w'}}{\partial s}(0, 1)$ is continuous, there exists a constant $C > 0$ (that only depends on r and on the geometry of (M, g)) such that

$$\|D_s \frac{\partial \Gamma_{w, w'}}{\partial s}(0, 1)\| \leq C,$$

for all $u, v \in T_{x_0}M$ with $\|u\|_{x_0} \leq r$ and $\|v\|_{x_0} \leq 1$. Hence, (6) yields that $\phi''(0) \geq -LC \|v\|_{x_0}^2$. Therefore, by setting $\alpha_1 = LC$, we have established that $F + \frac{\alpha_1}{2} \|\cdot\|_{x_0}^2$ is convex on $\{u \in T_{x_0}M : \|u\|_{x_0} \leq r\}$.

Step 2: Case where f is no longer assumed to be smooth.

Step 2.1: Smooth approximation of f . For all positive integers n , let $f_n : M \rightarrow \mathbb{R}$ be the function defined by

$$f_n(x) = n^d \int_{T_x M} f(\text{Exp}_x(v)) k(n\|v\|_x) d\Omega_x(v)$$

where $k : \mathbb{R} \rightarrow \mathbb{R}$ is a non-negative, smooth function supported on $[0, 1]$ with $\int_{\mathbb{R}} k(t) dt = 1$ and Ω_x is the pushforward measure of the Riemannian volume onto $T_x M$ by the exponential map Exp_x . By (Greene and Wu, 1973, Theorem 2), each f_n is smooth and the sequence $(f_n)_{n \geq 1}$ converges to f uniformly on K and satisfies

$$\liminf_{n \rightarrow \infty} \inf_{x \in K} \lambda_{\min}(\mathbf{H}f_n(x)) \geq 0,$$

where λ_{\min} stands for the smallest eigenvalue. Fix $\varepsilon > 0$, that we will then choose to be small enough. Without loss of generality, assume that $\inf_{x \in K} \lambda_{\min}(\mathbf{H}f_n(x)) \geq -\varepsilon$, for all $n \geq 1$ (since one might as well forget about the first terms of the sequence $(f_n)_{n \geq 1}$).

Step 2.2: f_n is Lipschitz on $B(x_0, r/2)$. Let us show that f_n is Lipschitz on $B(x_0, r/2)$, so long as n is large enough. Let $x, y \in B(x_0, r/2)$ and let $\pi_{x,y} : T_x M \rightarrow T_y M$ be the parallel transport map from x to y . Since $\pi_{x,y}$ is an isometry, it maps the measure Ω_x to Ω_y , hence, it holds that

$$f_n(y) = n^d \int_{T_x M} f(\text{Exp}_y(\pi_{x,y}(v))) k(n\|v\|_x) d\Omega_x(v).$$

Thus,

$$\begin{aligned} |f_n(x) - f_n(y)| &= n^d \left| \int_{T_x M} (f(\text{Exp}_x(v)) - f(\text{Exp}_y(\pi_{x,y}(v)))) k(n\|v\|_x) d\Omega_x(v) \right| \\ &\leq n^d \int_{T_x M} |f(\text{Exp}_x(v)) - f(\text{Exp}_y(\pi_{x,y}(v)))| k(n\|v\|_x) d\Omega_x(v) \\ &\leq Ln^d \int_{T_x M} d(\text{Exp}_x(v), \text{Exp}_y(\pi_{x,y}(v))) k(n\|v\|_x) d\Omega_x(v) \end{aligned} \quad (7)$$

where the last inequality holds as soon as n is large enough (i.e., $n \geq 2/r$) since f is L -Lipschitz on $K = B(x_0, r)$ and in the last integral, only the vectors $v \in T_x M$ with $\|v\|_{x_0} \leq n^{-1} \leq r/2$ contribute to the integral. Finally, the map $(x, y, v) \mapsto \text{Exp}_y(\pi_{x,y}(v))$ is smooth, hence it is Lipschitz on any compact set, so there exists $C > 0$ that is independent of x and y such that for all $v \in T_x M$ with $\|v\|_{x_0} \leq r/2$, $d(\text{Exp}_x(v), \text{Exp}_y(\pi_{x,y}(v))) \leq Cd(x, y)$. Therefore, (7) yields

$$|f_n(x) - f_n(y)| \leq LCd(x, y)n^d \int_{T_x M} k(n\|v\|_x) d\Omega_x(v) = LCd(x, y), \quad (8)$$

for all $x, y \in B(x_0, r/2)$. Thus, f_n is CL -Lipschitz on $B(x_0, r/2)$, for all integers $n \geq 2/r$.

Step 2.3: Add a fixed function to each f_n to make them convex. Assume that r and ε are chosen according to Lemma 12 below and let ϕ be the function given in that lemma. Then, for all $n \geq 1$, $f_n + \phi$ is convex on K , since its Hessian is positive semi-definite at any $x \in K$.

Step 2.4: Apply the smooth case (Step 1) to each $f_n + \phi$ on $B(x_0, r/2)$. For all $n \geq 1$, $f_n + \phi$ is smooth. Moreover, f_n is CL -Lipschitz on $B(x_0, r/2)$, where $C > 0$ does not depend on f , and ϕ is smooth so it is L' -Lipschitz on $B(x_0, r/2)$, for some $L' > 0$ that does not depend on f either. Therefore, applying the result proven in Step 1, we obtain that there exists $\alpha_2 > 0$ that does not depend on f , such that for all large enough integers n , $f_n \circ \text{Exp}_{x_0} + \phi \circ \text{Exp}_{x_0} + \frac{\alpha_2}{2} \|\cdot\|_{x_0}^2$ is convex on $\{u \in T_{x_0}M : \|u\|_{x_0} \leq r/2\}$. By taking the limit as $n \rightarrow \infty$, we obtain that $f \circ \text{Exp}_{x_0} + \phi \circ \text{Exp}_{x_0} + \frac{\alpha_2}{2} \|\cdot\|_{x_0}^2$ is convex on $\{u \in T_{x_0}M : \|u\|_{x_0} \leq r/2\}$.

Since ϕ is smooth, so is $\phi \circ \text{Exp}_{x_0}$, so there exists $\beta > 0$ such that $\lambda_{\max}(\bar{H}(\phi \circ \text{Exp}_{x_0})(u)) \leq \beta$, for all $u \in T_{x_0}M$ with $\|u\|_{x_0} \leq r/2$. Therefore, by setting $\alpha_3 = \alpha_2 + \beta$, we obtain that $f \circ \text{Exp}_{x_0} + \frac{\alpha_3}{2} \|\cdot\|_{x_0}^2$ is convex on $\{u \in T_{x_0}M : \|u\|_{x_0} \leq r/2\}$.

Step 3: Conclusion. We can now combine all cases by taking $\alpha = \max(\alpha_1, \alpha_3)$. □

Lemma 12

Let (M, g) be a complete Riemannian manifold without boundary and $x_0 \in M$. There exists $r, \varepsilon > 0$ and a smooth function $\phi : M \rightarrow \mathbb{R}$ such that $\lambda_{\min}(H\phi(x)) \geq \varepsilon$ for all $x \in B(x_0, r)$.

Proof

Let r_1 such that $B(x_0, r_1)$ is simply connected. For all $x \in M$ and for all linearly independent vectors $u, v \in T_xM$ with $\|u\|_x = \|v\|_x = 1$, let $\kappa(x, u, v)$ be the sectional curvature of M at x , along the 2-plane spanned by u and v (see (Do Carmo, 1992, Proposition 3.1 and Definition 3.2)). Then, κ is smooth, hence, there exists $\kappa_0 < \infty$ that is an upper bound on all sectional curvatures of M at points $x \in B(x_0, r_1)$, which is compact. Therefore, Lemma 22 yields that $B(x_0, r_1)$ is $\text{CAT}(\kappa_0)$. The desired result follows by taking: $r = \min(r_1, D_{\kappa_0}/4)$ and $\phi = \frac{\varepsilon}{2\alpha_{2r, \kappa_0}} d(\cdot, x_0)^2$, where D_{κ_0} and α_{2r, κ_0} are defined in Lemma 21, see Appendix D.2. ■

Acknowledgments

This work was funded by the French Agence Nationale de la Recherche, as part of the project ANR-22-ERCS-0015-01.

References

- Eddie Aamari and Clément Levrard. Nonasymptotic rates for manifold, tangent space and curvature estimation. 2019.
- Martial Agueh and Guillaume Carlier. Barycenters in the wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011.
- Adil Ahidar-Coutrix, Thibaut Le Gouic, and Quentin Paris. Convergence rates for empirical barycenters in metric spaces: curvature, convexity and extendable geodesics. *Probability theory and related fields*, 177(1-2):323–368, 2020.

- Stephanie Alexander, Vitali Kapovitch, and Anton Petrunin. Alexandrov geometry: preliminary version no. 1. *arXiv e-prints*, pages arXiv–1903, 2019.
- Jason M. Altschuler and Enric Boix-Adsera. Wasserstein barycenters can be computed in polynomial time in fixed dimension. *J. Mach. Learn. Res.*, 22:44–1, 2021.
- Jason M. Altschuler and Enric Boix-Adsera. Wasserstein barycenters are np-hard to compute. *SIAM Journal on Mathematics of Data Science*, 4(1):179–203, 2022.
- Kimon Antonakopoulos, Elena Veronica Belmega, and Panayotis Mertikopoulos. Online and stochastic optimization beyond lipschitz continuity: A riemannian approach. In *ICLR 2020-International Conference on Learning Representations*, pages 1–20, 2020.
- Rajendra Bhatia. Positive definite matrices. In *Positive Definite Matrices*. Princeton university press, 2009.
- Rajendra Bhatia and John Holbrook. Riemannian geometry and matrix geometric means. *Linear algebra and its applications*, 413(2-3):594–618, 2006.
- Rajendra Bhatia, Tanvi Jain, and Yongdo Lim. On the Bures-Wasserstein distance between positive definite matrices. *Expositiones Mathematicae*, 37(2):165–191, 2019.
- Rabi Bhattacharya and Lizhen Lin. Omnibus CLTs for Fréchet means and nonparametric inference on non-euclidean spaces. *Proceedings of the American Mathematical Society*, 145(1):413–428, 2017.
- Rabi Bhattacharya and Vic Patrangenaru. Large sample theory of intrinsic and extrinsic sample means on manifolds. *Annals of statistics*, 31(1):1–29, 2003.
- Rabi Bhattacharya and Vic Patrangenaru. Large sample theory of intrinsic and extrinsic sample means on manifolds: II. *Annals of statistics*, pages 1225–1259, 2005.
- Louis J. Billera, Susan P. Holmes, and Karen Vogtmann. Geometry of the space of phylogenetic trees. *Advances in Applied Mathematics*, 27(4):733–767, 2001.
- Martin R Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319. Springer Science & Business Media, 2013.
- Dmitri Burago, Yuri Burago, and Sergei Ivanov. *A course in metric geometry*, volume 33. American Mathematical Society, 2001.
- Yuan Shih Chow and Henry Teicher. *Probability theory: independence, interchangeability, martingales*. Springer Science & Business Media, 2003.
- Sebastian Clatici, Edward Chien, and Justin Solomon. Stochastic Wasserstein barycenters. In *International Conference on Machine Learning*, pages 999–1008. PMLR, 2018.
- Christopher Criscitiello and Nicolas Boumal. An accelerated first-order method for non-convex optimization on manifolds. *Foundations of Computational Mathematics*, pages 1–77, 2022.

- Marco Cuturi and Arnaud Doucet. Fast computation of Wasserstein barycenters. In *International conference on machine learning*, pages 685–693. PMLR, 2014.
- Manfredo Perdigao Do Carmo. *Riemannian geometry*, volume 6. Springer, 1992.
- Benjamin Eltzner and Stephan F Huckemann. A smeary central limit theorem for manifolds with application to high-dimensional spheres. *The Annals of Statistics*, 47(6):3360–3381, 2019.
- Benjamin Eltzner, Fernando Galaz-Garcia, Septhan F. Huckemann, and Wilderich Tuschmann. Stability of the cut locus and a central limit theorem for Fréchet means of Riemannian manifolds. *arXiv preprint arXiv:1909.00410*, 2019.
- Maurice Fréchet. Les éléments aléatoires de nature quelconque dans un espace distancié. *Ann. Inst. H. Poincaré*, 10:215–310, 1948. ISSN 0365-320X.
- Kei Funano. Rate of convergence of stochastic processes with values in \mathbb{R} -trees and Hadamard manifolds. *Osaka J. Math.*, 47(4):911–920, 2010.
- Christopher R Genovese, Marco Perone Pacifico, Verdinelli Isabella, Larry Wasserman, et al. Minimax manifold estimation. *Journal of machine learning research*, 13:1263–1291, 2012.
- Robert E. Greene and Hung Wu. On the subharmonicity and plurisubharmonicity of geodesically convex functions. *Indiana University Mathematics Journal*, 22(7):641–653, 1973.
- Shelby J. Haberman. Concavity and estimation. *The Annals of Statistics*, pages 1631–1661, 1989.
- CH Himmelberg, T Parthasarathy, and FS Van Vleck. On measurable relations. *Fundamenta mathematicae*, 61:161–167, 1982.
- Thomas Hotz, Sean Skwerer, Stephan Huckemann, Huiling Le, J Stephen Marron, Jonathan C Mattingly, Ezra Miller, James Nolen, Megan Owen, and Vic Patrangenaru. Sticky central limit theorems on open books. *Annals of Applied Probability*, 2013.
- Peter J Huber. Robust estimation of a location parameter. *Breakthroughs in statistics: Methodology and distribution*, pages 492–518, 1992.
- Stephan F. Huckemann. Intrinsic inference on the mean geodesic of planar shapes and tree discrimination by leaf growth. *The Annals of Statistics*, 39(2):1098–1124, 2011.
- Hermann Karcher. Riemannian center of mass and mollifier smoothing. *Communications on pure and applied mathematics*, 30(5):509–541, 1977.
- David George Kendall, Dennis Barden, Thomas K. Carne, and Huiling Le. *Shape and shape theory*, volume 500. John Wiley & Sons, 2009.
- Wilfrid S. Kendall and Huiling Le. Limit theorems for empirical Fréchet means of independent and non-identically distributed manifold-valued random variables. *Brazilian Journal of Probability and Statistics*, 25(3):323 – 352, 2011.
- Alexey Kroshnin, Nazarii Tupitsa, Darina Dvinskikh, Pavel Dvurechensky, Alexander Gasnikov, and Cesar Uribe. On the complexity of approximating Wasserstein barycenters. In *International conference on machine learning*, pages 3530–3540. PMLR, 2019.

- Alexey Kroshnin, Vladimir Spokoiny, and Alexandra Suvorikova. Statistical inference for Bures-Wasserstein barycenters. *The Annals of Applied Probability*, 31(3):1264–1298, 2021.
- Thibaut Le Gouic and Jean-Michel Loubes. Existence and consistency of Wasserstein barycenters. *Probability Theory and Related Fields*, 168(3):901–917, 2017.
- Thibaut Le Gouic, Quentin Paris, Philippe Rigollet, and Austin J Stromme. Fast convergence of empirical barycenters in Alexandrov spaces and the Wasserstein space. *Journal of the European Mathematical Society*, 2022.
- John M. Lee. *Introduction to smooth manifolds*. Springer, 2012.
- John M. Lee. *Introduction to Riemannian manifolds*, volume 2. Springer, 2018.
- Yongdo Lim and Miklós Pálfi. Weighted deterministic walks for the least squares mean on Hadamard spaces. *Bulletin of the London Mathematical Society*, 46, 05 2014.
- Estelle Massart, Julien M. Hendrickx, and P-A Absil. Curvature of the manifold of fixed-rank positive-semidefinite matrices endowed with the Bures-Wasserstein metric. In *Geometric Science of Information: 4th International Conference, GSI 2019, Toulouse, France, August 27–29, 2019, Proceedings*, pages 739–748. Springer, 2019.
- Wojciech Niemiro. Asymptotics for M-estimators defined by convex minimization. *The Annals of Statistics*, pages 1514–1533, 1992.
- Shin-ichi Ohta. Convexities of metric spaces. *Geom. Dedicata*, 125:225–250, 2007. ISSN 0046-5755.
- Shin-ichi Ohta and Miklós Pálfi. Discrete-time gradient flows and law of large numbers in Alexandrov spaces. *Calc. Var. Partial Differential Equations*, 54(2):1591–1610, 2015.
- R Tyrrell Rockafellar. *Convex analysis*, volume 18. Princeton university press, 1970.
- Rolf Schneider. *Convex bodies: the Brunn–Minkowski theory*. Number 151. Cambridge university press, 2014.
- Christof Schötz. Convergence rates for the generalized fréchet mean via the quadruple inequality. *Electronic Journal of Statistics*, 13:4280–4345, 2019.
- Nicolas Schreuder, Victor-Emmanuel Brunel, and Arnak Dalalyan. Statistical guarantees for generative models without domination. In *Algorithmic Learning Theory*, pages 1051–1071. PMLR, 2021.
- Karl-Theodor Sturm. Probability measures on metric spaces of nonpositive curvature. In *Heat kernels and analysis on manifolds, graphs, and metric spaces (Paris, 2002)*, volume 338 of *Contemp. Math.*, pages 357–390. Amer. Math. Soc., Providence, RI, 2003.
- Constantin Udriste. *Convex functions and optimization methods on Riemannian manifolds*, volume 297. Springer Science & Business Media, 2013.
- Aad W Van Der Vaart and Jon A Wellner. *Weak convergence*. Springer, 1996.

Daniel H. Wagner. Survey of measurable selection theorems. *SIAM Journal on Control and Optimization*, 15(5):859–903, 1977.

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

Hongyi Zhang and Suvrit Sra. Towards riemannian accelerated gradient methods. *arXiv preprint arXiv:1806.02812*, 2018.

Herbert Ziezold. On expected figures and a strong law of large numbers for random elements in quasi-metric spaces. In *Transactions of the Seventh Prague Conference on Information Theory, Statistical Decision Functions, Random Processes and of the 1974 European Meeting of Statisticians*, pages 591–602. Springer, 1977.

Appendix A. On measurable selections of subgradients

Lemma 13 *Let $(\mathcal{Z}, \mathcal{F}, P)$ be a probability space and (M, g) be a Riemannian manifold without boundary. Let $\phi : \mathcal{Z} \times M \rightarrow \mathbb{R}$ be a function such that*

- $\phi(z, \cdot)$ is convex for P -almost all $z \in \mathcal{Z}$;
- $\phi(\cdot, x)$ is measurable for all $x \in M$.

Then, there exists a map $g : \mathcal{Z} \times M \rightarrow TM$ such that

- $g(z, x) \in T_x M$ is a subgradient of $\phi(z, \cdot)$ at x , for all $x \in M$, for P -almost all $z \in \mathcal{Z}$;
- $g(\cdot, x)$ is measurable, for all $x \in M$.

Proof

Let $N \in \mathcal{F}$ be such that $P(N) = 0$ and $\phi(z, \cdot)$ is convex on M for all $z \in \mathcal{Z} \setminus N$. Set $\tilde{\mathcal{Z}} = \mathcal{Z} \setminus N$. It is enough to prove that for all $x \in M$, there exists a measurable function $g(\cdot, x)$ on M such that $g(z, x)$ is a subgradient of $\phi(z, \cdot)$ at x , for all $z \in \tilde{\mathcal{Z}}$.

Fix $x_0 \in M$. For all $z \in \mathcal{Z}$, let $\Gamma(z)$ be the set of all subgradients of $\phi(z, \cdot)$ at the point x_0 . By (Udriste, 2013, Theorem 4.5), $\Gamma(z) \neq \emptyset$, for all $z \in \tilde{\mathcal{Z}}$. Therefore, it is enough to show that the restriction of Γ to $\tilde{\mathcal{Z}}$ is weakly measurable, i.e., that $\{z \in \tilde{\mathcal{Z}} : \Gamma(z) \cap U \neq \emptyset\} \in \mathcal{F}$, for all open subsets $U \subseteq T_{x_0} M$ (see (Himmelberg et al., 1982) for the definition).

One can write $\Gamma(z)$ as

$$\begin{aligned} \Gamma(z) &= \bigcap_{v \in T_{x_0} M} \{u \in T_{x_0} M : \phi(z, \text{Exp}_{x_0}(v)) \geq \phi(z, x_0) + \langle u, v \rangle_{x_0}\} \\ &= \bigcap_{v \in E_0} \{u \in T_{x_0} M : \phi(z, \text{Exp}_{x_0}(v)) \geq \phi(z, x_0) + \langle u, v \rangle_{x_0}\}, \end{aligned}$$

where E_0 is some fixed countable, dense subset of $T_{x_0} M$. The last equality follows from the continuity of $\phi(z, \cdot)$, by Lemma 6. For all $v \in E_0$, let $\Gamma_v(z) = \{u \in T_{x_0} M : \phi(z, \text{Exp}_{x_0}(v)) \geq \phi(z, x_0) + \langle u, v \rangle_{x_0}\}$.

Since each of the countably many $\Gamma_v(z)$, $v \in E_0$, is closed, (Himmelberg et al., 1982, Corollary 4.2) shows that it is enough to show that each Γ_v is weakly measurable. Fix $v \in E_0$ and let U be an open subset of $T_{x_0} M$. Then, for all $z \in \tilde{\mathcal{Z}}$,

$$\begin{aligned} \Gamma_v(z) \cap U \neq \emptyset &\iff \exists u \in U, \phi(z, \text{Exp}_{x_0}(v)) \geq \phi(z, x_0) + \langle u, v \rangle_{x_0} \\ &\iff \exists u \in U \cap E_0, \phi(z, \text{Exp}_{x_0}(v)) \geq \phi(z, x_0) + \langle u, v \rangle_{x_0}. \end{aligned}$$

The last equivalence follows from the fact that since U is open, $U \cap E_0$ is dense in U , and if some $u \in U$ satisfies $\phi(z, \text{Exp}_{x_0}(v)) \geq \phi(z, x_0) + \langle u, v \rangle_{x_0}$, then, any $u' \in U$ of the form $u' = u - v'$, for any $v' \in T_{x_0} M$ with small enough norm such that $\langle v, v' \rangle_{x_0} \leq 0$. There is at least one such u' that falls in $U \cap E_0$.

Finally, one obtains:

$$\{z \in \tilde{\mathcal{Z}} : \Gamma_v(z) \cap U \neq \emptyset\} = \bigcup_{u \in E_0} \{z \in \tilde{\mathcal{Z}} : \phi(z, \text{Exp}_{x_0}(v)) \geq \phi(z, x_0) + \langle u, v \rangle_{x_0}\},$$

which is in \mathcal{F} since $\phi(\cdot, x)$ is assumed to be measurable for all $x \in M$.

One concludes using Lemma 14 below. ■

Lemma 14

Let $(\mathcal{Z}, \mathcal{F}, P)$ be a probability space and E a Polish space. Let Γ be a function that maps any $z \in \mathcal{Z}$ to a closed subset of E . Assume the existence of $N \in \mathcal{F}$ with $P(N) = 0$ and such that $\Gamma(z) \neq \emptyset$ for all $z \in \mathcal{Z} \setminus N$. Assume further that for all open subsets U of E , $\{z \in \mathcal{Z} \setminus N : \Gamma(z) \cap U \neq \emptyset\} \in \mathcal{F}$. Then, there exists a function $g : \mathcal{Z} \rightarrow E$ such that $g(z) \in \Gamma(z)$, for P -almost all $z \in \mathcal{Z}$.

Proof

Consider the probability space $(\tilde{\mathcal{Z}}, \tilde{\mathcal{F}}, \tilde{P})$ defined by setting

- $\tilde{\mathcal{Z}} = \mathcal{Z} \setminus N$;
- $\tilde{\mathcal{F}} = \{F \setminus N : F \in \mathcal{F}\}$;
- \tilde{P} as the restriction of P to $(\tilde{\mathcal{Z}}, \tilde{\mathcal{F}})$.

Then, for all $z \in \mathcal{Z}$, $\Gamma(z) \neq \emptyset$ so, thanks to (Wagner, 1977, Theorem 4.1), one can find a measurable selection g of the restriction of Γ to $\tilde{\mathcal{Z}}$. By setting $g(z)$ to be any constant value in E for $z \in N$, we obtain a function $g : \mathcal{Z} \rightarrow E$ that satisfies the requirement. ■

Appendix B. Elements on Riemannian manifolds and on geodesic convexity

In this section, let (M, g) be a complete d -dimensional Riemannian manifold.. That is, M is a d -dimensional smooth manifold and the Riemannian metric g equips each tangent space $T_x M$, $x \in M$, with a dot product $g_x(u, v)$, $u, v \in T_x M$. Moreover, by Hopf-Rinow theorem (Do Carmo, 1992, Chapter 7, Theorem 2.8), completeness of (M, d) ensures that any two points $x, y \in M$ are connected by at least minimizing geodesic, i.e., a smooth path $\gamma : [0, 1] \rightarrow M$ such that $\gamma(0) = x$, $\gamma(1) = y$ and $d(\gamma(s), \gamma(t)) = |s - t|d(x, y)$, for all $s, t \in [0, 1]$.

B.1. On smooth functions and their derivatives

If $f : M \rightarrow \mathbb{R}$ is a smooth function, its gradient is the map ∇f such that for all $x \in M$, $\nabla f(x) \in T_x M$ and for all $u \in T_x M$,

$$\langle \nabla f(x), u \rangle_x = (f \circ \gamma)'(0)$$

where $\gamma : [0, 1] \rightarrow M$ is any smooth path such that $\gamma'(0) = u$. The Hessian of f at any $x \in M$ is represented by the linear map $\mathbf{H}f(x) : T_x M \rightarrow T_x M$ such that for all $u \in T_x M$,

$$\langle u, \mathbf{H}f(x)u \rangle_x = (f \circ \gamma)''(0)$$

where γ is a geodesic path with $\gamma(0) = x$, $\gamma'(0) = u$.

With our notation, any smooth function $f : M \rightarrow \mathbb{R}$ has a second order Taylor expansion at any given point $x \in M$ of the following form:

$$f(\text{Exp}_x(tu)) = f(x) + t\langle \nabla f(x), u \rangle_x + \frac{t^2}{2}\langle u, \mathbf{H}f(x)u \rangle_x + o(t^2)$$

as $t \rightarrow 0$, for any fixed $u \in T_x M$.

B.2. On convex functions

Recall that a convex function on M is any function $f : M \rightarrow \mathbb{R}$ that satisfies $f(\gamma(t)) \leq (1-t)f(x) + tf(y)$ for all $t \in [0, 1]$, $x, y \in M$ and for all constant speed, minimizing geodesic $\gamma : [0, 1] \rightarrow M$ from x to y . Just as in Euclidean spaces, convex functions are automatically continuous (Udriste, 2013, Theorem 3.6) (note that we only consider convex functions on the whole space M , which is assumed to have no boundary). The notion of subgradients can also be extended to the Riemannian case.

Definition 15 *Let $f : M \rightarrow \mathbb{R}$ be a convex function and $x \in M$. A vector $v \in T_x M$ is called a subgradient of M at x if and only if*

$$f(\text{Exp}_x(u)) \geq f(x) + \langle v, u \rangle_x,$$

for all $u \in T_x(M)$.

The set of all subgradients of f at $x \in M$ is denoted by $\partial f(x)$ and (Udriste, 2013, Theorem 4.5) ensures that $\partial f(x) \neq \emptyset$, for all $x \in M$ (again, note that we only consider the case when M has no boundary). For simplicity here, unlike in (Udriste, 2013), subgradients are elements of the tangent space, not its dual.

Subgradients of convex functions satisfy the following monotonicity property.

Lemma 16

Let $f : M \rightarrow \mathbb{R}$ be a convex function. Let $x \in M$, $u \in T_x M$ and set $y = \text{Exp}_x(u)$. Let $g_1 \in \partial f(x)$ and $g_2 \in \partial f(y)$. Then,

$$\langle u, \pi_{x,y}^{-1}(g_2) - g_1 \rangle_x \geq 0,$$

where $\pi_{x,y}$ is the parallel transport from x to y through the geodesic given by $\gamma(t) = \text{Exp}_x(tu)$, $t \in [0, 1]$.

Proof

First, note that by letting $v = \pi_{x,y}(u)$, it holds that $x = \text{Exp}_y(-v)$. By definition of subgradients, one has the following inequalities:

$$f(y) \geq f(x) + \langle u, g_1 \rangle_x$$

and

$$f(x) \geq f(y) + \langle -v, g_2 \rangle_y = f(y) - \langle u, \pi_{x,y}^{-1}(g_2) \rangle_x,$$

using the fact that $\pi_{x,y}$ is an isometry. Summing both inequalities yields the result. ■

Appendix C. On sequences of convex functions in metric spaces

C.1. Proof of Lemma 4

Almost sure convergence Since (M, d) is a proper metric space, it is separable (Alexander et al., 2019, Proposition 2.2.2). Let M_0 be a countable dense subset of M . By assumption, $F_n(x) \xrightarrow[n \rightarrow \infty]{} F(x)$ almost surely, for all $x \in M_0$. Therefore, with probability one, the convergence holds for all $x \in M_0$, since M_0 is countable. Hence, because M_0 is also dense, Theorem 3 applies and we obtain the desired result.

Convergence in probability Here, we use the following fact.

Lemma 17 (Chow and Teicher, 2003, Lemma 2 (ii) p.67) *Let X, X_1, X_2, \dots be real valued random variables. Then, X_n converges in probability to X if and only if every subsequence of $(X_n)_{n \geq 1}$ has itself a subsequence that converges almost surely to X .*

Let K be a compact subset of M and consider a subsequence of $(\sup_{x \in K} |F_n(x) - F(x)|)_{n \geq 1}$. Without loss of generality (one could renumber that subsequence), we actually consider the whole sequence $(\sup_{x \in K} |F_n(x) - F(x)|)_{n \geq 1}$ and we show that it admits a subsequence that converges almost surely to zero. Let M_0 be a countable dense subset of M . Denote the elements of M_0 as x_1, x_2, \dots

By assumption, $F_n(x_1)$ converges to $F(x_1)$ in probability, therefore, it has a subsequence, say $F_{\phi_1(n)}(x_1)$, that converges to $F(x_1)$ almost surely.

Now, again by assumption, $F_{\phi_1(n)}(x_2)$ converges to $F(x_2)$ in probability, therefore, it has a subsequence, say $F_{\phi_1 \circ \phi_2(n)}(x_2)$, that converges to $F(x_2)$ almost surely.

By iterating this procedure, we build, for all integers $k \geq 1$, an increasing function $\phi_k : \mathbb{N}^* \rightarrow \mathbb{N}^*$ such that $F_{\phi_1 \circ \dots \circ \phi_k(n)}(x_k)$ converges almost surely to $F(x_k)$, as $n \rightarrow \infty$.

Finally, let $\psi : \mathbb{N}^* \rightarrow \mathbb{N}^*$ the function defined by $\psi(n) = \phi_1 \circ \dots \circ \phi_n(n)$, for all $n \geq 1$. Then, by construction, $F_{\psi(n)}(x_k)$ converges almost surely to $F(x_k)$, for all $k \geq 1$. Hence, with probability 1, $F_{\psi(n)}(x)$ converges almost surely to $F(x)$ for all $x \in M_0$. Lemma 3 yields that $\sup_{x \in K} |F_{\psi(n)}(x) - F(x)|$ converges to zero almost surely, which concludes the proof.

Appendix D. A brief introduction to Alexandrov's curvature

In this section, we give the precise definition of CAT spaces.

D.1. Model spaces

First, we introduce a family of model spaces that will allow us to define local and global curvature bounds in the sequel. Let $\kappa \in \mathbb{R}$.

$\kappa = 0$: **Euclidean plane** Set $M_0 = \mathbb{R}^2$, equipped with its Euclidean metric. This model space corresponds to zero curvature, is a geodesic space where geodesics are unique and given by line segments.

$\kappa > 0$: **Sphere** Set $M_\kappa = \frac{1}{\sqrt{\kappa}}\mathbb{S}^2$: This is the 2-dimensional Euclidean sphere, embedded in \mathbb{R}^3 , with center 0 and radius $1/\sqrt{\kappa}$, equipped with the arc length metric: $d_\kappa(x, y) = \frac{1}{\sqrt{\kappa}} \arccos(\kappa x^\top y)$, for all $x, y \in M_\kappa$. This is a geodesic space where the geodesics are unique except for opposite

points, and given by arcs of great circles. Here, a great circle is the intersection of the sphere with any plane going through the origin.

$\kappa < 0$: **Hyperbolic space** Set $M_\kappa = \frac{1}{\sqrt{\kappa}}\mathbb{H}^2$, where $\mathbb{H}^2 = \{(x_1, x_2, x_3) \in \mathbb{R}^3 : x_3 > 0, x_1^2 + x_2^2 - x_3^2 = -1\}$. The metric is given by $d_\kappa = \frac{1}{\sqrt{-\kappa}}\operatorname{arccosh}(-\kappa\langle x, y \rangle)$, for all $x, y \in M_\kappa$, where $\langle x, y \rangle = x_1y_1 + x_2y_2 - x_3y_3$. This is a geodesic space where geodesics are always unique and are given by arcs of the intersections of M_κ with planes going through the origin.

For $\kappa \in \mathbb{R}$, let D_κ be the diameter of the model space M_κ , i.e., $D_\kappa = \begin{cases} \infty & \text{if } \kappa \leq 0 \\ \frac{\pi}{\sqrt{\kappa}} & \text{if } \kappa > 0 \end{cases}$.

D.2. Curvature bounds

Let (M, d) be a geodesic space. The notion of curvature (lower or upper) bounds for (M, d) is defined by comparing the triangles in M with triangles with the same side lengths in model spaces.

Definition 18 *A (geodesic) triangle in M is a set of three points in M (the vertices) together with three geodesic segments connecting them (the sides).*

Given three points $x, y, z \in S$, we abusively denote by $\Delta(x, y, z)$ a triangle with vertices x, y, z , with no mention to which geodesic segments are chosen for the sides (geodesics between points are not necessarily unique, as seen for example on a sphere, between any two opposite points). The perimeter of a triangle $\Delta = \Delta(x, y, z)$ is defined as $\operatorname{per}(\Delta) = d(x, y) + d(y, z) + d(x, z)$. It does not depend on the choice of the sides.

Definition 19 *Let $\kappa \in \mathbb{R}$ and Δ be a triangle in M with $\operatorname{per}(\Delta) < 2D_\kappa$. A comparison triangle for Δ in the model space M_κ is a triangle $\bar{\Delta} \subseteq M_\kappa$ with same side lengths as Δ , i.e., if $\Delta = \Delta(x, y, z)$, then $\bar{\Delta} = \Delta(\bar{x}, \bar{y}, \bar{z})$ where $\bar{x}, \bar{y}, \bar{z}$ are any points in M_κ satisfying*

$$\begin{cases} d(x, y) = d_\kappa(\bar{x}, \bar{y}) \\ d(y, z) = d_\kappa(\bar{y}, \bar{z}) \\ d(x, z) = d_\kappa(\bar{x}, \bar{z}). \end{cases}$$

Note that a comparison triangle in M_κ is always unique up to rigid motions. We are now ready to define curvature bounds. Intuitively, we say that (M, d) has global curvature bounded from above (resp. below) by κ if all its triangles (with perimeter smaller than $2D_\kappa$) are thinner (resp. fatter) than their comparison triangles in the model space M_κ .

Definition 20 *Let $\kappa \in \mathbb{R}$. We say that (M, d) has global curvature bounded from above (resp. below) by κ if and only if for all triangles $\Delta \subseteq M$ with $\operatorname{per}(\Delta) < 2D_\kappa$ and for all $x, y \in \Delta$, $d(x, y) \leq d_\kappa(\bar{x}, \bar{y})$ (resp. $d(x, y) \geq d_\kappa(\bar{x}, \bar{y})$), where \bar{x} and \bar{y} are the points on a comparison triangle $\bar{\Delta}$ in M_κ that correspond to x and y . We then call (M, d) a $\operatorname{CAT}(\kappa)$ (resp. $\operatorname{CAT}^+(\kappa)$) space.*

Lemma 21 (*Ohta, 2007, Proposition 3.1*)

Let (M, d) be $CAT(\kappa)$ for some $\kappa \in \mathbb{R}$ and assume that M has diameter $D < D_\kappa/2$. Then, for all $x_0 \in M$, the function $d(\cdot, x_0)$ is convex on M . Moreover, set $\alpha_{D,\kappa} = \frac{\sqrt{\kappa}D}{\tan(\sqrt{\kappa}D)}$ if $\kappa > 0$, $\alpha_{D,\kappa} = 1$ otherwise. Then, the function $\frac{1}{2}d^2(\cdot, x_0)$ is $\alpha_{D,\kappa}$ -strongly convex, i.e., it satisfies

$$\frac{1}{2}d^2(\gamma(t), x_0) \leq \frac{1-t}{2}d^2(x, x_0) + \frac{t}{2}d^2(y, x_0) - \frac{\alpha_{D,\kappa}t(1-t)}{4}d(x, y)^2,$$

for all $x, y \in M$ and $\gamma \in \Gamma_{x,y}$.

D.3. Examples of CAT spaces

In this section, we give a few examples of CAT spaces. First, the following lemma, which is a direct consequence of (Bridson and Haefliger, 2013, Theorem 1A.6) together with (Burago et al., 2001, Theorem 9.2.9), give a whole class of Riemannian manifolds as CAT spaces.

Lemma 22 *Let (M, g) be a simply connected Riemannian manifold. Assume that all the sectional curvatures are bounded from above by some $\kappa \in \mathbb{R}$. Then, M is $CAT(\kappa)$.*

We give two examples that are of particular relevance in matrix analysis and optimal transport.

1. Let M be the collection of all $p \times p$ real, symmetric positive definite matrices. Equip M with the metric given by $d(A, B) = \|\log(A^{-1/2}BA^{-1/2})\|_F$, where \log is the matrix logarithm and $\|\cdot\|_F$ is the Frobenius norm. Then, d is inherited from a Riemannian metric and (M, d) is $CAT(0)$ (Bhatia and Holbrook, 2006). This metric space is particularly important in the study of matrix geometric means. For any $A, B \in M$, their geometric mean is defined as $A\#B = A^{1/2}(A^{-1/2}BA^{-1/2})^{1/2}A^{1/2}$, and it is the unique midpoint from A to B , i.e., $A\#B = \gamma(1/2)$, where γ is the unique geodesic between A and B .
2. Let M be the collection of all $p \times p$ real, symmetric positive definite matrices with spectrum included in $[\lambda_0, \infty)$, for some $\lambda_0 > 0$. Equip M with the Bures-Wasserstein metric, obtained as follows. For $A, B \in M$, let $d(A, B) = W_2(\mathcal{N}_p(0, A), \mathcal{N}_p(0, B))$, where W_2 is the Wasserstein distance and $\mathcal{N}_p(0, A)$ (resp. $\mathcal{N}_p(0, B)$) is the p -variate Gaussian distribution with mean 0 and covariance matrix A (resp. B). Then, d is also inherited from a Riemannian metric (Bhatia et al., 2019) and (M, d) is $CAT(3/(2\lambda_0))$, by (Massart et al., 2019, Proposition 2).

Further examples are given below.

- Metric trees are $CAT(\kappa)$ for any $\kappa \in \mathbb{R}$ (since all its triangles are flat) (Bridson and Haefliger, 2013, Section II.1.15)
- The surface of a smooth convex body in a Euclidean space is $CAT(\kappa)$, where $\kappa > 0$ is an upper bound on the reach of the complement of the body, as a consequence of (Schneider, 2014, Theorem 3.2.9)
- The space of phylogenetic trees (Billera et al., 2001) is $CAT(0)$.