

Conformal-DP: Data Density Aware Privacy on Riemannian Manifolds via Conformal Transformation

Peilin He[†]

ORCID: 0000-0003-3553-9949

Department of Informatics and Networked Systems
University of Pittsburgh
Pittsburgh, Pennsylvania, USA
peilin.he@pitt.edu

M. Amin Rahimian

ORCID: 0000-0001-9384-1041

Department of Industrial Engineering
University of Pittsburgh
Pittsburgh, Pennsylvania, USA
rahimian@pitt.edu

Liou Tang

ORCID: 0009-0005-5220-5176

Department of Informatics and Networked Systems
University of Pittsburgh
Pittsburgh, Pennsylvania, USA
liou.tang@pitt.edu

James Joshi[†]

ORCID: 0000-0003-4519-9802

Department of Informatics and Networked Systems
University of Pittsburgh
Pittsburgh, Pennsylvania, USA
jjoshi@pitt.edu

[†]Corresponding author.

Abstract—Differential Privacy (DP) enables privacy-preserving data analysis by adding calibrated noise. While recent works extend DP to curved manifolds (e.g., diffusion-tensor MRI, social networks) by adding geodesic noise, these assume uniform data distribution. This assumption is not always practical, hence these approaches may introduce biased noise and suboptimal privacy-utility trade-offs for non-uniform data. To address this issue, we propose *Conformal-DP* that utilizes conformal transformations on Riemannian manifolds. This approach locally equalizes sample density and redefines geodesic distances while preserving intrinsic manifold geometry. Our theoretical analysis demonstrates that the conformal factor, which is derived from local kernel density estimates, is data density-aware. We show that under these conformal metrics, *Conformal-DP* satisfies ϵ -differential privacy on any complete Riemannian manifold and offers a closed-form expected geodesic error bound dependent only on the maximal density ratio, and not global curvature. We show through experiments on synthetic and real-world datasets that our mechanism achieves superior privacy-utility trade-offs, particularly for heterogeneous manifold data, and also is beneficial for homogeneous datasets.

I. INTRODUCTION

With the exponential growth of data containing significant privacy-sensitive information, we are increasingly facing the critical challenge of data sharing for analytics or training AI models in a privacy-preserving manner. These challenges are further exacerbated when such data reside in non-Euclidean spaces, such as manifolds or graph-structured domains, where traditional linear algebraic methods fall short, e.g., medical imaging data such as MRI or CT scans [1]–[3] often represent anatomical structures that conform to curved surfaces or volu-

metric manifolds. Similarly, flat Euclidean metrics cannot fully capture geographical data such as terrain elevation models or climate patterns on a spherical Earth [4], [5]. In addition, such data are challenging for model training and inference in the computer vision area [6]–[8] which are prone to generate only locally effective yet globally suboptimal results. Therefore, these datasets call for more advanced modeling techniques, such as geometric- or manifold-based methods. These data may contain sensitive information; thus, enabling safe sharing of such data with formal privacy guarantees is necessary.

Differential Privacy (DP), first introduced by Dwork et al. in [9], has emerged as a gold standard for data privacy, in particular, for sharing sensitive datasets for analysis, because of its rigorous mathematical foundations in guaranteeing privacy by limiting the risk of re-identification of individuals/samples in a dataset. Several variations and enhancements of DP mechanisms have been proposed in the literature, from standard DP [9]–[11] to Rényi-DP [12], Concentrated-DP [13] and Random-DP [14], that allow generating differentially private datasets suitable for privacy-preserving data analytics or AI model development. DP mechanisms involve adding perturbations to transform the original data to protect privacy. Most of these DP mechanisms mainly focus on linear or Euclidean data and are not suitable for non-linear data [15]–[17].

Existing works by Reimherr et al. in [18] and Utpala et al. in [19] have extended Laplacian and Gaussian DP to Riemannian manifolds, respectively. The former demonstrates that it is possible to design DP mechanisms strictly relying on

intrinsic geometric distances and volumes defined naturally on the manifolds, while the latter shows a significant improvement of Gaussian DP on manifolds compared to Laplacian perturbations. However, these mechanisms rely on a uniform privacy perturbation across the manifold \mathcal{M} regardless of the data density. We argue that this hampers utility of the DP mechanism when there is a heterogeneous data distribution, while not offering significant levels of privacy preservation. Consider a dataset with heterogeneous data density across \mathcal{M} , where dense and sparse regions of the data would intrinsically require different levels of perturbation to guarantee their contribution to the dataset and privacy needs, as discussed in existing works [20]–[23].

In this paper, we propose a novel data-density-aware differential privacy mechanism to address the above research questions by leveraging local data density to control the addition of noise. An overview of the mechanism is in Figure 1. Our key contributions are as follows:

- We propose a Conformal DP mechanism that applies different levels of perturbations to different data subsets by accounting for their *data density*.
- We utilize *Conformal Transformation* [24], [25], [25], [26], [26], [27], an intrinsic geometric tool in Riemannian geometry, to encode data density in the underlying metric to guide the privacy perturbation.
- We construct a novel smoothed *conformal factor* that reshapes the local geometry of the manifold, ensuring that high data density regions are compressed while low data density regions are expanded.
- We present theoretical results to ensure soundness of the mechanism by showing formally a bound on privacy loss using bi-Lipschitz continuity between the original and conformal geodesic distances.
- We demonstrate experimental results to show that our data density-aware approach yields better privacy-utility tradeoffs for manifold data, and simultaneously unifies conventional DP methods with manifold-valued extensions.

The remainder of this paper is organized as follows. Section II provides an overview of foundational concepts and background for differential privacy (DP), Riemannian geometry, and conformal transformations, while introducing key notations used throughout the work in Table I. Section III presents the details for constructing of the *conformal factor* on the Riemannian manifold \mathcal{M} via conformal transformation and related conformal metrics. Building on the conformal factor, Section IV introduces our proposed *Conformal-DP* mechanism, a novel privacy-preserving method that leverages geometric properties of the manifold. Section V demonstrates the theoretical analysis of the privacy-utility trade-offs inherent in our proposed mechanism, quantifying its optimality under bounded curvature constraints. Section VI and VII develops the algorithms and reports on experimental results that validate the theoretical results using both synthetic and real-world datasets, compare our approach with existing techniques, and

TABLE I: Key notations used in this paper.

Original Riemannian Manifolds Notations	
\mathcal{M}	A compact Riemannian manifold.
d	Dimension of manifold \mathcal{M} .
g	Original Riemannian metric on \mathcal{M} .
$\langle \cdot, \cdot \rangle_m$	Riemannian metric at point m on \mathcal{M} .
$\gamma(t)$	A smooth path or geodesic connecting points on \mathcal{M} .
$\ \dot{\gamma}(t)\ _{\gamma(t)}$	Norm of the velocity vector of γ at point $\gamma(t)$.
$\rho_g(x, y)$	Geodesic distance between points x, y under original metric g .
$T_m\mathcal{M}$	Tangent space at m under g .
$\text{inj}(\mathcal{M})$	Injectivity radius of the manifold \mathcal{M} .
$\mu_g, d\mu_g$	Volume measure associated with / induced by metric g .
g_{ij}	Components of the Riemannian metric tensor g in local coordinates.
$B_r(m)$	Geodesic ball with central point m with radius r under metric g .
Δ_g	Laplace–Beltrami (LB) operator associated with metric g .
$g_{\mu\nu}$	Metric tensor in a d -dimensional space.
$\eta_{\mu\nu}$	Flat metric (Euclidean or Minkowski) in conformal transformations.
$\varepsilon^\mu(x)$	Infinitesimal parameter representing a small conformal transformation.
Conformal Metric Notations	
g^*	Conformal metric defined as $g^* = e^{2\sigma}g$.
$\rho_{g^*}(x, y)$	Geodesic distance between points x, y under conformal metric g^* .
μ_{g^*}	Volume measure associated with conformal metric g^* .
$\sigma(x)$	Conformal scaling function, solved from PDEs.
$\phi(x)$	Conformal factor defined as $\phi(x) = e^{2\sigma(x)}$.
ϕ_{\min}, ϕ_{\max}	Lower and upper bounds of conformal factor $\phi(x)$.
λ^*	Rate parameter of the conformal Laplace mechanism.
Differential Privacy Mechanism Notations	
Δ	Global sensitivity under original metric g .
$\Delta^*(D)$	Local sensitivity under conformal metric g^* given dataset D .
$f_{\text{data}}(x)$	Kernel density function on \mathcal{M} .
ε	Privacy budget parameter
$\mathcal{A}(\cdot)$	A random mechanism.
$L_{D, D'}(z)$	Privacy loss random variable comparing datasets D and D' .
$\mathbb{P}_r^*(\cdot \cdot)$	Probability Density Function with respect to the measure μ_{g^*}
$\eta(D)$	Output summary (Fréchet mean) of dataset D on \mathcal{M} .
$\eta_\rho(D)$	Privatized output summary (Fréchet mean) of dataset D on \mathcal{M} .

show the mechanism’s practicality in balancing privacy guarantees with utility preservation. Finally, we conclude the paper with a discussion of use cases, limitations, and future research directions.

II. PRELIMINARIES AND BACKGROUND

In this section, we present basic notations and an overview of key concepts related to Differential Privacy, Riemannian manifolds, and conformal transformation. For more details, we refer the readers to Dwork et al. [9], [11], [28] for DP, and to Reimherr et al. and Jiang et al. [18], [29] for extensions of DP over Riemannian manifolds. Further, for more details on the Riemannian manifolds and the conformal space transformation, there are several works such as [24], [27], [30], [31] that provide theoretical background.

A. Differential Privacy

Differential privacy (DP) is a rigorous statistical framework designed to protect individual data entries while allowing meaningful analysis of sensitive datasets. Let \mathcal{X} be the domain of all possible user records, and let $D \in \mathcal{X}^n$ be a dataset of n records. A dataset $D' \in \mathcal{X}^n$ is called an adjacent (or neighboring) dataset of D if D and D' differ in at most one record. A randomized mechanism M , which takes a dataset D as input and outputs a result in some range \mathbb{R} , is said to be ε -differentially private (ε -DP) if, for every pair of adjacent datasets (D, D') that differ in only one element and every measurable subset $S \subseteq \mathbb{R}$, the following holds

$$\Pr[M(D) \in S] \leq e^\varepsilon \Pr[M(D') \in S]. \quad (1)$$

Here, the privacy loss is quantified by the privacy budget ε : the smaller the value of ε , the better the privacy achieved. The probabilities are taken over the randomness of the mechanism M , which typically injects random noise or employs sampling to limit the impact of any single record on the output. In practice, analyses often involve multiple differentially private computations.

In addition, geometrical optimization for DP mechanisms has also been explored. We build on a line of work on Differential Privacy over Riemannian manifolds [32]–[35]. Reimherr et al. in [18] first propose ε -DP on Riemannian manifolds; these efforts highlight the promise of Riemannian DP in contexts where the data naturally lie on curved spaces. By leveraging the intrinsic geometry of the manifold, one can inject noise while respecting local curvature and preserving important geometric structures. We can generalize classical DP mechanisms to curved domains, e.g., Laplace DP mechanisms [18], [36] or Gaussian DP [29], [37] mechanisms.

B. Riemannian Geometry

Here, we provide an overview of relevant background on Riemannian geometry [2], [32]–[35].

Riemannian Manifolds. Throughout the paper, we use \mathcal{M} to represent a d -dimensional complete Riemannian manifold. For each point $m \in \mathcal{M}$, let $T_m\mathcal{M}$ be the tangent space at m . The manifold \mathcal{M} is equipped with a smoothly varying *Riemannian metric*:

$$\{\langle \cdot, \cdot \rangle_m : m \in \mathcal{M}\}, \quad (2)$$

Which provides an inner product on $T_m\mathcal{M}$ for every $m \in \mathcal{M}$.

Geodesics and Riemannian Distance. Given two points $m_1, m_2 \in \mathcal{M}$, consider a smooth path $\gamma : [0, 1] \rightarrow \mathcal{M}$ with endpoints $\gamma(0) = m_1$ and $\gamma(1) = m_2$. The *velocity vector* $\dot{\gamma}(t)$ of this path lies in the tangent space $T_{\gamma(t)}\mathcal{M}$. The *length* of γ is defined by

$$L(\gamma) = \int_0^1 \|\dot{\gamma}(t)\|_{\gamma(t)} dt = \int_0^1 \left(\langle \dot{\gamma}(t), \dot{\gamma}(t) \rangle_{\gamma(t)} \right)^{\frac{1}{2}} dt. \quad (3)$$

The *Riemannian distance* ρ between m_1 and m_2 is then given by

$$\rho(m_1, m_2) := \inf_{\substack{\gamma: \gamma(0)=m_1 \\ \gamma(1)=m_2}} L(\gamma). \quad (4)$$

Any path γ that realizes this infimum is called a *geodesic*. A geodesic is the natural generalization of a straight line in Euclidean space for a curved Riemannian manifold.

Exponential Map and Injectivity Radius. A key tool for relating the manifold \mathcal{M} to its tangent spaces is the *exponential map*, $\exp_m : T_m\mathcal{M} \rightarrow \mathcal{M}$. If a unique geodesic γ from m_1 to m_2 exists, then $\exp_{m_1}(\dot{\gamma}(0)) = m_2$. Under completeness property, \exp_m is locally a diffeomorphism at each m . Its local inverse, called the *logarithm map* \exp_m^{-1} ,

is well-defined in a suitable neighborhood of m . The largest radius of such a neighborhood is called the *injectivity radius* at m . The *injectivity radius* of \mathcal{M} , denoted $\text{inj}(\mathcal{M})$, is the infimum of these radii over all points $m \in \mathcal{M}$.

Riemannian Volume Measure. The Riemannian metric on \mathcal{M} induces a natural *volume measure* μ . In local coordinates given by a chart:

$$\varphi : U \subset \mathcal{M} \rightarrow \varphi(U) \subset \mathbb{R}^d, \quad (5)$$

With components:

$$g_{ij} = \langle \partial_i, \partial_j \rangle_m, \quad (6)$$

Where $\{\partial_i\}$ are the coordinate basis vectors on $T_m\mathcal{M}$. In an n -dimensional Riemannian manifold, one has the canonical volume form $d\mu_g$. In local coordinates x^1, x^2, \dots, x^n , where:

$$g = g_{ij} dx^i \otimes dx^j, \quad (7)$$

And the volume form is

$$d\mu_g = \sqrt{\det(g_{ij})} dx^1 \wedge \dots \wedge dx^n. \quad (8)$$

where \wedge is the wedge product from exterior calculus. Globally, this form is well defined (up to orientation) using a partition of unity¹, which defines the volume measure μ on \mathcal{M} .

C. Conformal Transformation over Riemannian Manifolds

A conformal transformation (conformal diffeomorphism) is a smooth map $f : (\mathcal{M}, g) \rightarrow (\mathcal{M}^*, g^*)$ that preserves angles; equivalently, it rescales the metric by a positive function. In particular,

$$g^* = e^{2\sigma} g, \quad \sigma : \mathcal{M} \rightarrow \mathbb{R},$$

where σ is the *conformal factor* [24], [27], [30]. When $\mathcal{M}^* = \mathcal{M}$, f is a conformal symmetry of (\mathcal{M}, g) [38].

The angle-preserving condition implies that lengths are stretched or contracted by $e^{2\sigma(x)}$, yet the inner product is uniformly scaled, so all angles between tangent vectors (and thus between smooth curves) remain unchanged [39]. If $e^{2\sigma}$ is constant, f is homothetic, scaling all lengths uniformly [40]. Thus, conformal diffeomorphisms keep the qualitative shape of infinitesimal figures intact while altering distances; they map the geometry of \mathcal{M} to that of \mathcal{M}^* up to a position-dependent scale factor. Metrics linked by $g^* = e^{2\sigma} g$ lie in the same conformal class [24].

We adopt $g^* = e^{2\sigma} g$ because (i) σ is obtained from the elliptic PDE of Section III, guaranteeing existence, uniqueness, and smoothness, and (ii) This multiplicative form gives explicit bi-Lipschitz bounds on distances and volumes, crucial for the privacy-utility analysis in Section IV.

¹A *partition of unity* in \mathcal{M} that is a collection of smooth non-negative functions $\{\varphi_i\}$ whose sum is identically 1 in \mathcal{M} . Each function φ_i has support contained in a coordinate chart, allowing local data to be smoothly extended to the whole manifold. For details, see [35, Chapter 2].

III. CONSTRUCTING DATA DENSITY AWARE CONFORMAL METRICS

In this section, we construct *conformal factor* ϕ and associate metrics (*conformal geodesic distance*, *conformal volume metric*) on the Riemannian manifold \mathcal{M} . These constructs serve as the foundational components for our subsequent DP mechanism design. The remainder of this section is organized as follows. In Section III-A, we establish the existence and uniqueness of the conformal scaling function $\sigma(x)$ by solving the Poisson equation $\Delta_g \sigma(x) = -f_{\text{data}}(x)$. Using $\sigma(x)$, we then define a bounded, density-aware conformal factor $\phi(x) = e^{2\sigma(x)}$ and the conformal metric $g^* = e^{2\sigma}g$. Subsequently, Sections III-B and III-C analyze the relationship between the original and conformal metrics, proving that the conformal transformation preserves bi-Lipschitz bounds and consistency for probability measures.

A. Construction of the Conformal Factor

Let \mathcal{M} be a d -dimensional complete Riemannian manifold equipped with Riemannian metric $g : \{(\cdot, \cdot)_m : m \in \mathcal{M}\}$ and Borel σ -algebra [41], [42]. To define $\sigma(x)$ on \mathcal{M} , we adopt a generalized approach with the Laplace–Beltrami (LB) operator Δ_g . Specifically, we consider $\sigma(x)$ as a solution to a Poisson-type PDE:

$$\Delta_g \sigma(x) = -f_{\text{data}},$$

where f_{data} encodes the underlying data density via KDE [43]; $\Delta_g = \text{div}_g \nabla$ [44]. This PDE naturally ties $\sigma(x)$ to the geometric properties of \mathcal{M} to achieve the desired curvature under a conformal transformation. Assume \mathcal{M} is compact without boundary [45]. Under these conditions, standard elliptic regularity theory ensures that if $f_{\text{data}}(x)$ is smooth, then $\sigma(x)$ is also smooth [46], [47].

Define the conformal factor $\phi(x) = e^{2\sigma(x)}$. It works as follows: In denser regions, this factor is smaller, contracting distances. In sparser regions, the factor is larger (up to a maximum), stretching distances. This conformal adjustment provides varying utility based on data density, introducing greater metric distortion in low-density, potentially more sensitive areas. The following two lemmas give the prerequisites for defining the conformal factor.

Lemma III.1 ([48], [49]). *Let (\mathcal{M}, g) be a closed Riemannian manifold. Denote a smooth function $H \in C^\infty(\mathcal{M})$. If $\Delta_g \sigma = H$ with $\sigma \in C^\infty(\mathcal{M})$ and mean-zero H , then*

$$\int_{\mathcal{M}} H \, d\mu_g = \int_{\mathcal{M}} \Delta_g \sigma \, d\mu_g = 0.$$

Conversely, if $\int_{\mathcal{M}} H \, d\mu_g = 0$, the Poisson equation $\Delta_g \sigma = H$ admits a smooth solution on compact manifolds [50].

Lemma III.2 ([51]). *Under the same setting, any two solutions σ_1, σ_2 satisfy $\Delta_g(\sigma_1 - \sigma_2) = 0$; hence $\sigma_1 - \sigma_2$ is harmonic.*

By the maximum principle on a compact manifold, a harmonic function is constant, so solutions are unique up to

an additive constant. Fixing $\int_{\mathcal{M}} \sigma \, d\mu_g = 0$ yields a canonical solution.

$$H = -(f_{\text{data}} - c),$$

where $c = \frac{1}{\text{vol}(\mathcal{M})} \int_{\mathcal{M}} f_{\text{data}} \, d\mu_g$. Lemma III.1 guarantees a smooth solution σ , and Lemma III.2 ensures uniqueness. Set

$$\phi(x) = e^{2\sigma(x)} > 0, \quad g^* = \phi g.$$

Elliptic regularity implies $\sigma, \phi \in C^\infty$, so g^* is a smooth Riemannian metric as well.

Mirshani *et al.* [52] observed that sensitivity depends on the noise distribution; our *conformal-DP* mechanism rescales \mathcal{M} and applies geometry-consistent noise.² The preceding lemmas motivate the following theorem, which demonstrates that a general second-order, strongly elliptic operator admits a smooth solution for the conformal scaling function σ . This solution produces a strictly positive conformal factor $\phi = e^{2\sigma}$, ensuring that the rescaled metric $g^* = \phi \cdot g$ is well defined and smooth.

Theorem III.3. *Let (\mathcal{M}, g) be a smooth Riemannian manifold and Ω_g a second-order, strongly elliptic operator with sufficiently smooth coefficients. Let $H \in C^\infty(\mathcal{M} \times \mathbb{R})$ be a prescribed smooth nonlinear function. Then the nonlinear equation*

$$\Omega_g[\sigma](x) = H(x, \sigma(x)), \quad \forall x \in \mathcal{M},$$

has a smooth solution $\sigma \in C^\infty(\mathcal{M})$ (the conformal scaling function). Consequently, the conformal factor

$$\phi(x) := e^{2\sigma(x)} > 0$$

is smooth, and the conformal metric

$$g^*(x) := e^{2\sigma(x)} g(x) = \phi(x) g(x)$$

is also smooth and positive-definite on \mathcal{M} .

Proof. We present a detailed proof in Appendix A. \square

Boundaries of the Conformal Factor: To prevent the conformal factor function $\phi(x)$ from being infinite or prevent it from degenerating on the manifold \mathcal{M} , further to produce extreme conformal metric g^* , we propose Theorem III.4 to give the upper and lower bounds for the conformal factor $\phi(x)$.

Theorem III.4 (Upper and Lower Bounds). *Let Theorem III.3, Lemma III.1 and Lemma III.2 hold, denote $\lambda_1 > 0$ be the smallest nonzero eigenvalue of $-\Delta_g$ (i.e., its spectral gap [53]). Define $f_{\min} = \min_{x \in \mathcal{M}} f_{\text{data}}(x)$, $f_{\max} = \max_{x \in \mathcal{M}} f_{\text{data}}(x)$. Then the following bounds hold:*

1) *Operator norm bound:*

$$\|\sigma\|_{L^\infty(\mathcal{M})} \leq \frac{1}{\lambda_1} \|f_{\text{data}}\|_{L^\infty(\mathcal{M})}.$$

2) *Extremal bounds:*

$$-\frac{f_{\max}}{\lambda_1} \leq \sigma(x) \leq -\frac{f_{\min}}{\lambda_1} \text{ for all } x \in \mathcal{M}.$$

²There is rarely an explicit form for Δ_g ; we discretize Δ_g and H over $D \subset \mathcal{M}^n$ to solve for σ (see and Table II).

Proof. We present a detailed proof in Appendix B. \square

Theorem III.4 guarantees the following bounds for the conformal factor $\phi(x)$: $e^{2(-f_{\max}/\lambda_1)} \leq \phi(x) \leq e^{2(-f_{\min}/\lambda_1)}$.

B. Conformal Geodesic Distance

After constructing the conformal factor, we start to obtain other key components. Consider the local coordinates given by $\{x_1, \dots, x_n\}$. By the Hopf-Rinow theorem [35, Theorem 6.19], this means that for every pair of points, there exists a geodesic that minimizes the distance between them [18]. For a differentiable curve $\gamma : [a, b] \rightarrow \mathcal{M}$, such that $\gamma(t) = (\gamma^1(t), \dots, \gamma^n(t))$, the length of this curve under the conformal metric g^* is given by: $L_{g^*}(\gamma) = \int_a^b \sqrt{g_{\gamma(t)}^*(\dot{\gamma}(t), \dot{\gamma}(t))} dt$. Note that $g^*(x) = \phi(x)g(x)$, it follows that $g_{\gamma(t)}^*(\dot{\gamma}(t), \dot{\gamma}(t)) = \phi(\gamma(t))g_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t))$. Therefore, we have the following.

$$L_{g^*}(\gamma) = \int_a^b \sqrt{\phi(\gamma(t))} \sqrt{g_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t))} dt. \quad (9)$$

Consider points $x, y \in \mathcal{M}$, the geodesic distance under the conformal metric g^* is as follows:

$$\rho_{g^*}(x, y) = \inf_{\substack{\gamma: [a, b] \rightarrow \mathcal{M} \\ \gamma(a)=x, \gamma(b)=y}} L_{g^*}(\gamma), \quad (10)$$

This means among all differential curves γ from x to y , we take the minimum value of $L_{g^*}(\gamma)$. Combining Eq. (9) and Eq. (10), we have:

$$\rho_{g^*}(x, y) = \inf_{\gamma(a)=x, \gamma(b)=y} \int_a^b \sqrt{\phi(\gamma(t))} \sqrt{g_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t))} dt. \quad (11)$$

It is important to note that, in general, the geodesic distance in the conformal metric cannot be expressed in a simple closed form using $\rho_g(x, y)$ because the geodesics in g^* differ from those in g unless the conformal factor ϕ is constant. Since ϕ is continuous on the compact manifold \mathcal{M} , it attains a minimum $\phi_{\min} > 0$ and a maximum ϕ_{\max} (extreme-value theorem [54]). Based on the Uniform Boundedness Principle provided by the following lemma in existing work:

Lemma III.5 ([55]). *There exist constants $\phi_{\min} = \inf_{x \in \mathcal{M}} \phi(x) > 0$ and $\phi_{\max} = \sup_{x \in \mathcal{M}} \phi(x) > 0$, such that for every $x \in \mathcal{M}$, the inequality $\phi_{\min} \leq \phi(x) \leq \phi_{\max}$ holds uniformly.*

We then can compare the lengths of arbitrary curves $\gamma : [0, 1] \rightarrow \mathcal{M}$ under original and conformal metrics: (\mathcal{M}, g) and (\mathcal{M}, g^*) . Based on Eq. (9), if $t \in [0, 1]$, $\sqrt{\phi_{\min}} \leq \sqrt{\phi(\gamma(t))} \leq \sqrt{\phi_{\max}}$, we present the Corollary III.6 based on Lemma III.5 and Eq. (11) and bi-Lipschitz comparison [56]:

Corollary III.6. *For a complete Riemannian manifold (\mathcal{M}, g) and a conformal metric $g^* = \phi(x)g(x)$, if $\phi(x) \in [\phi_{\min}, \phi_{\max}]$, for all $(x, y) \in \mathcal{M}$, there exists*

$$\sqrt{\phi_{\min}} L_g(\gamma) \leq L_{g^*}(\gamma) \leq \sqrt{\phi_{\max}} L_g(\gamma)$$

, where $L_g(\gamma) = \int_a^b \sqrt{g_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t))} dt$. Thus, for every pair of arbitrary points in $(x, y) \in \mathcal{M}$, the geodesic distance under g^* and the geodesic distance under g satisfy the following bilateral estimates:

$$\sqrt{\phi_{\min}} \rho_g(x, y) \leq \rho_{g^*}(x, y) \leq \sqrt{\phi_{\max}} \rho_g(x, y), \text{ for all } x, y \in \mathcal{M}. \quad (12)$$

Proof. We present a detailed proof in Appendix C. \square

The above inequality relationships avoid explicitly solving for the geodesic distance under the conformal metric (because when the conformal factor σ is nonconstant, geodesics are often difficult to find), but provide useful global distance estimates. This will be used as key inequality equations related to sensitivity or distance to analyze differential privacy.

C. Conformal Volume Transformation

Similar to geodesic distance, the Riemannian volume μ changes under a conformal transformation [57]. Schoen et al. [25] and Topping et al. [58] study how conformal changes in the metric affect volume elements and curvature. In local coordinates $\{x^1, \dots, x^n\}$, the Riemannian volume is

$$d\mu_g(x) = \sqrt{\det(g_{ij}(x))} dx^1 \wedge \dots \wedge dx^n,$$

where $[g_{ij}(x)]_{i,j=1}^n$ is the metric matrix. Under the conformal transformation of Section II-C, the new metric matrix is

$$[g_{ij}^*(x)]_{i,j=1}^n = \phi(x) [g_{ij}(x)]_{i,j=1}^n,$$

so that

$$\det(g_{ij}^*(x)) = \det(\phi(x) g_{ij}(x)) = \phi(x)^n \det(g_{ij}(x)).$$

Hence the conformal volume is represent as:

$$\begin{aligned} d\mu_{g^*}(x) &= \sqrt{\det(g_{ij}^*(x))} dx^1 \wedge \dots \wedge dx^n \\ &= \sqrt{\phi(x)^n \det(g_{ij}(x))} dx^1 \wedge \dots \wedge dx^n \\ &= \phi(x)^{\frac{n}{2}} \sqrt{\det(g_{ij}(x))} dx^1 \wedge \dots \wedge dx^n \\ &= \phi(x)^{\frac{n}{2}} d\mu_g(x). \end{aligned} \quad (13)$$

The next section presents our proposed *Conformal-DP* mechanism based on conformal factor, conformal geodesic distance, and conformal volumes.

IV. CONFORMAL-DP MECHANISM

Existing works have demonstrated that achieving ε -DP is feasible in measure spaces equipped with the Borel σ -algebra [41], [42]. Building on this theoretical foundation, we construct a Laplacian-type DP mechanism by incorporating the conformal factor ϕ into the geometry of the underlying space. We use ϕ to induce the conformal transformation of the base metric, dynamically modulating the spatial distribution of the noise.

Because ϕ is smooth and varies continuously on the compact manifold, the output of the perturbation mechanism remains aligned with the manifold's shape at every point. We give an overview of our *Conformal-DP* mechanism in Fig.1.

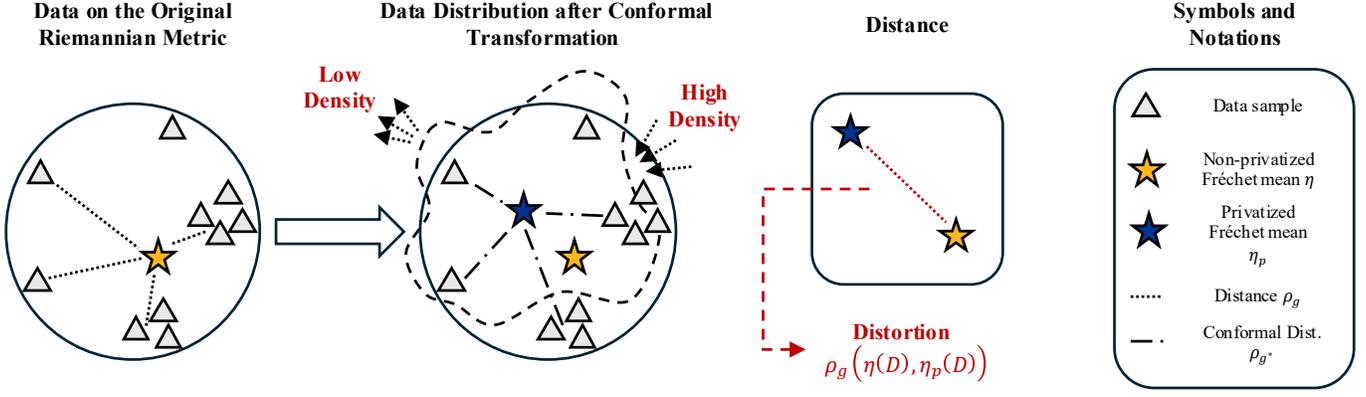


Fig. 1: An overview of our proposed Conformal-DP mechanism.

A. Topological Analysis

Let \mathcal{M} be a compact and smooth Riemannian manifold of dimension n and it is equipped with: 1) Riemannian metric g ; 2) Data-Density-Aware conformal metric $g^* = \phi \cdot g$, where the conformal factor $\phi : \mathcal{M} \rightarrow \mathbb{R} > 0$. The Riemannian manifolds (\mathcal{M}, g) and (\mathcal{M}, g^*) share the same underlying smooth topological structure. Specifically, the identity map $\text{id}_{\mathcal{M}} : (\mathcal{M}, g) \rightarrow (\mathcal{M}, g^*)$ is a diffeomorphism [24]–[26], ensuring that any point z in the topological manifold \mathcal{M} is identified with itself, independently of g or g^* for all $z \in \mathcal{M}$, so that the topology and smooth structure of \mathcal{M} remain unchanged. The metrics g and g^* differ only in their geometric measurements (e.g., geodesic distances and volumes) on \mathcal{M} . For example, consider a stochastic output $\vec{m} \in \mathcal{M}$ calculated under the metric g^* . Because \vec{m} is fundamentally a point in the topological manifold \mathcal{M} , no additional “mapping” is required to interpret \vec{m} in (\mathcal{M}, g) . The difference arises only in the way the geometric properties of \vec{m} are quantified under g versus g^* .

B. Distribution Analysis

Building on the topological foundations of Section IV-A, consider a dataset $D = \{x_1, \dots, x_n\}$ defined on the compact Riemannian manifold (\mathcal{M}, g) . Let f be a statistical summary for the dataset on \mathcal{M} , where $D \subseteq B_r(\eta)$ and

$$B_r(\eta) = \{x \in \mathcal{M} \mid \rho_g(\eta, x) < r\},$$

with ρ_g denoting geodesic distance under metric g . Thus, the same logic happens under the conformal metric $g^* = \phi \cdot g$, the geodesic ball becomes

$$B_r^*(\eta) = \{x \in \mathcal{M} \mid \rho_{g^*}(\eta, x) < r\},$$

where ρ_{g^*} is the conformal geodesic distance; these balls inherently have finite volume.

However, we do not calculate $f(D)$ under the conformal metric, since g^* is used solely for noise distribution rather than redefining the statistic itself. The non-privatized summary remains $f(D)$ on (\mathcal{M}, g) . We use ρ_{g^*} and μ_{g^*} to control the

spatial distribution of noise and obtain a randomly perturbed f_p . This simplifies control over normalization constants for distributions supported in $B_r^*(\eta)$. Specifically,

$$\int_{B_r^*(\eta)} \exp[-\lambda^* \rho_{g^*}(\eta, x)] d\mu_{g^*}(x)$$

remains finite without additional curvature constraints. Consequently, on compact manifolds, one can localize noise within $B_r^*(\eta)$ with exponential decay in $\rho_{g^*}(\eta, x)$ while maintaining proper normalization. To precisely articulate its impact on bounding the sensitivity of the statistical summary under data perturbations, we adopt Definition IV.1 from [18]:

Definition IV.1. (Reimherr et al. [18]) The data $D \subseteq B_r(m_0)$ for some m_0 , where $r < r^* := \frac{1}{2} \min\{\text{inj } \mathcal{M}, \frac{\pi}{2} \kappa^{-1/2}\}$ and $\kappa > 0$ is an upper bound on the sectional curvatures of \mathcal{M} , consider two datasets $D = \{x_1, \dots, x_{n-1}, x_n\}$ and $D' = \{x_1, \dots, x_{n-1}, x'_n\}$ differing by only one element. If \bar{x} and \bar{x}' are the two sample statistical summaries of D and D' respectively, then

$$\rho(\bar{x}, \bar{x}') \leq \frac{2r(2 - h(r, \kappa))}{nh(r, \kappa)}, \quad h(r, \kappa) = \begin{cases} 2r\sqrt{\kappa} \cot(\sqrt{\kappa}2r), & \kappa > 0, \\ 1, & \kappa \leq 0. \end{cases}$$

Therefore, $\Delta := \frac{2r(2 - h(r, \kappa))}{n, h(r, \kappa)}$ is an upper bound on the change in the statistical summary (Here we choose the most suitable metric Fréchet mean η) when a single data point is replaced. We now extend Definition IV.1 to the conformal metric. To utilize conformal geometric properties, we build an unnormalized Laplace-type kernel \tilde{K}_r^* based on heat kernel theorem [59] that is strictly positive inside the ball $B_r^*(\eta)$ of radius r centered at Fréchet mean η under the distance ρ_{g^*} and zero elsewhere:

$$\tilde{K}_r^*(z \mid \eta) = \exp\{-\lambda^* \rho_{g^*}(\eta, z)\} \mathbf{1}_{\{\rho_{g^*}(\eta, z) < r\}}, \quad (14)$$

where z is a variable in $B_r^*(\eta)$, and $\lambda^* > 0$ is an adjustable Laplace rate parameter under conformal metric g^* (we will define the noise scale parameter later). We refer to $\tilde{K}_r^*(\cdot \mid \eta)$ as an unnormalized kernel because $\int_{\mathcal{M}} \tilde{K}_r^*(z \mid \eta) d\mu_{g^*}(z)$ is in general finite, but not normalized to 1. In order to transform

\tilde{K}_r^* to a probability density for our proposed approach, we compute the normalization constant:

$$\begin{aligned} C_r(\eta, \lambda^*) &= \int_{\mathcal{M}} \tilde{K}_r^*(z | \eta) d\mu_{g^*}(z) \\ &= \int_{B_r^*(\eta)} \exp\{-\lambda^* \rho_{g^*}(\eta, z)\} d\mu_{g^*}(z), z \in B_r^*, \end{aligned} \quad (15)$$

where the integral is effectively restricted to the ball $B_r^*(\eta) = \{z : \rho_{g^*}(\eta, z) < r\}$. Next, we define the noise distribution on $B_r^*(\eta)$ as follows:

$$\mathbb{P}_r^*(z | \eta) = \frac{1}{C_r(\eta, \lambda^*)} \begin{cases} \exp\{-\lambda^* \rho_{g^*}(\eta, z)\}, & z \in B_r^*(\eta), \\ 0, & z \notin B_r^*(\eta). \end{cases} \quad (16)$$

Thus, $\mathbb{P}_r^*(z | \eta)$ is now a proper probability density function (pdf) with respect to the measure μ_{g^*} . Equivalently, it defines a probability measure \mathbb{P}_η on the original metric (\mathcal{M}, g) .

$$\mathbb{P}_\eta(S) = \int_S \mathbb{P}_r^*(z | \eta) d\mu_{g^*}(z), \quad \text{for all } S \subseteq \mathcal{M},$$

where S denotes a measurable set on \mathcal{M} . We refer to $\mathbb{P}_r^*(z | \eta)$ as the localized Laplace-type distribution centered at η under the conformal metric g^* . With this distribution, we present the following theorem that builds a geodesic-Laplace kernel with conformal metrics, and provides the foundation for the Conformal-DP mechanism \mathcal{A} :

Theorem IV.2. *Let $\rho_{g^*}(\cdot, \cdot)$ be the geodesic distance induced by g^* on \mathcal{M} , and μ_{g^*} be the associated Riemannian volume measure with respect to g^* . Consider a fixed point $\eta \in \mathcal{M}$ and radius $r > 0$, and $B_r^*(\eta) = \{x \in \mathcal{M} | \rho_{g^*}(\eta, x) < r\}$ be the open geodesic ball of radius r centered at η under the metric g^* . For a given rate parameter $\lambda^* > 0$,*

$$\tilde{K}_r^*(z | \eta) := \exp[-\lambda^* \rho_{g^*}(\eta, z)] \mathbf{1}_{\{\rho_{g^*}(\eta, z) < r\}}, z \in \mathcal{M}$$

be the unnormalized Laplace-type kernel, combining with (15), we obtain the probability density on \mathcal{M} with respect to the measure μ_{g^*} .

$$\begin{aligned} \mathbb{P}_r^*(z | \eta) &:= \frac{\tilde{K}_r^*(z | \eta)}{C_r(\eta, \lambda^*)} \\ &= \begin{cases} \frac{\exp[-\lambda^* \rho_{g^*}(\eta, z)]}{\int_{B_r^*(\eta)} \exp[-\lambda^* \rho_{g^*}(\eta, u)] d\mu_{g^*}(u)}, & z \in B_r^*(\eta), \\ 0, & z \notin B_r^*(\eta). \end{cases} \end{aligned} \quad (17)$$

Proof. We present a detailed proof in Appendix D. \square

By defining an exponentially decaying kernel supported only on a geodesic ball $B_r^*(\eta)$ and showing that its integral is finite, it establishes that normalizing this kernel yields a valid probability density. Thus, we obtain a density-aware localized noise distribution that decays with the conformal distance ρ_{g^*} around the Fréchet mean η .

C. Data-Density-Aware Perturbations

In this section, we propose the density-aware perturbation method under the conformal metric g^* . Recall from Section III-B, because $\phi(x)$ is determined by the dataset D , the conformal geodesic distance ρ_{g^*} will change depending on the local data density. In the traditional scenario [18], the global sensitivity is commonly defined as follows: if the output centroids of adjacent datasets D and D' are $\eta(D)$ and $\eta(D')$, then $\rho_g(\eta(D), \eta(D')) \leq \Delta$; here, we denote Δ as the global sensitivity of \mathcal{M} , (original metric in our statement). To formalize this sensitivity perspective under both the original and conformal metrics, we first recall the classical Laplace mechanism guarantee stated in the following definition:

Definition IV.3. ([18]) Consider two adjacent datasets D and D' , and let $f : \mathcal{X}^n \rightarrow \mathcal{M}$ be a summary with global sensitivity Δ . If $\rho(f(D), f(D')) \leq \Delta$, then the Laplace mechanism with footpoint $f(D)$ and rate $\sigma = 2\Delta/\varepsilon$ satisfies ε -differential privacy. If the normalizing constant, $C_{n, \sigma}$, does not depend on the footpoint, η , then one can take $\sigma = \Delta/\varepsilon$.

For conformal metrics g^* , we care about what the upper bound is on the central distance $\rho_{g^*}(\eta(D), \eta(D'))$ between adjacent datasets under g^* . Through Corollary III.6, we have:

$$\rho_{g^*}(\eta(D), \eta(D')) \leq \sqrt{\phi_{\max}} \rho_g(\eta(D), \eta(D')) = \sqrt{\phi_{\max}} \cdot \Delta \quad (18)$$

To this end, we extend the global sensitivity Δ on the original metric to the conformal global sensitivity Δ^* based on the density of data on \mathcal{M} ; that is, for each data set D , we define the sensitivity in the conformal metrics as follows:

$$\Delta^*(D) = \sup_{D' \in \mathcal{N}(D)} \rho_{g^*}(\eta(D), \eta(D')) \quad (19)$$

where $\mathcal{N}(D)$ represents the set of neighboring data sets of D . As a result of $\eta(D)$ possibly being located in different high or low density areas of the distribution, $\Delta^*(D)$ will also vary depending on the geometric properties of the data distribution. According to Eq.(26) for D and D' respectively: $\mathbb{P}^*(z | \eta(D), \lambda^*)$ and $\mathbb{P}^*(z | \eta(D'), \lambda^*)$, define the privacy loss random variable on the logarithmic scale [60]:

$$L_{D, D'}(z) = \ln \frac{\mathbb{P}^*(z | \eta(D), \lambda^*)}{\mathbb{P}^*(z | \eta(D'), \lambda^*)}. \quad (20)$$

Under the conditions of Theorem IV.2, we propose Theorem IV.4 as follows:

Theorem IV.4. *For the two neighboring datasets D and D' , suppose that the random mechanism \mathcal{A} produces a random output z in (\mathcal{M}, g^*) . We conclude that the mechanism \mathcal{A} is ε -DP if:*

$$L_{D, D'}(z) \leq \varepsilon \quad \text{and} \quad L_{D', D}(z) \leq \varepsilon, \quad \text{for all } z \in \mathcal{M},$$

and we obtain the following explicit form for $L_{D,D'}(z)$ by combining (17) and (20),

$$\begin{aligned} L_{D,D'}(z) &= -\lambda^* [\rho_{g^*}(\eta(D), z) - \rho_{g^*}(\eta(D'), z)] \\ &\quad + \ln \frac{\mathbb{P}(\eta(D'), \lambda^*)}{\mathbb{P}(\eta(D), \lambda^*)} \\ &\leq \varepsilon. \end{aligned} \quad (21)$$

Proof. We present a detailed proof in Appendix E. \square

In order to guarantee Eq.(21), we give Corollary IV.5 based on the privacy loss random variable:

Corollary IV.5. *Given two adjacent datasets D and D' , a randomized mechanism \mathcal{A} , the neighboring inputs $x = \eta(D) \sim x' = \eta(D')$ and every output z in \mathcal{M} , define the privacy loss function $\ell_{\mathcal{A},x,x'}(z)$ at outcome z as:*

$$\ell_{\mathcal{A},x,x'}(z) = \ln \frac{\mathbb{P}^*(z | x, \lambda^*)}{\mathbb{P}^*(z | x', \lambda^*)}. \quad (22)$$

Assume that $\mathbb{P}^*(\cdot | x, \lambda^*)$ and $\mathbb{P}^*(\cdot | x', \lambda^*)$ are absolutely continuous with respect to the same reference measure μ_{g^*} , with densities that are continuous and strictly positive everywhere. We claim that mechanism \mathcal{A} satisfies ε -differential privacy if and only if for every output $z \in \mathcal{M}$ and every pair of adjacent summaries $x \sim x'$, the privacy loss is bounded by ε in absolute value:

$$-\varepsilon \leq \ell_{\mathcal{A},x,x'}(z) = \ln \frac{\mathbb{P}^*(z | x, \lambda^*)}{\mathbb{P}^*(z | x', \lambda^*)} \leq \varepsilon \quad (23)$$

With detailed proof in Appendix F, it establishes an ε -DP guarantee for a random \mathcal{A} whose outputs lie in the conformal metric (\mathcal{M}, g^*) . The result makes it clear how to control ε -DP via the λ^* , the geometry of the space (through the metric g^* and the distance ρ_{g^*}). This allows users to see exactly how changes in the dataset translate into changes in the distribution over outcomes, and thus to control and certify privacy, the clear experiment result is shown in Figure 5.

Sensitivity Analysis. In Corollary III.6, we have proved that $\rho_{g^*}(\eta(D), \eta(D')) \leq \sqrt{\phi_{\max}} \cdot \Delta$, therefore, we use the upper bounds of the conformal factor to connect to the original manifold. Assume $x = f(D)$, $y = f(D')$; then from Section III-C we have the pre-requisites: for all $x, y \in \mathcal{M}$, Eq.(12) holds. In the original metric (\mathcal{M}, g) , we define a statistical summary $f(D) \in \mathcal{M}$, it takes dataset D and maps to a point in the manifold, here we use Fréchet mean [61] to represent $f(D)$. In the conformal metric (\mathcal{M}, g^*) , we want to obtain the conformal geodesic distance between their output points of comparison. Thus, if $\rho_g(x, y) \leq \Delta$, where Δ is the global sensitivity of \mathcal{M} , then we have the worst-case sensitivity: $\Delta^* = \rho_{g^*}(x, y) \leq \sqrt{\phi_{\max}} \cdot \Delta$ under the conformal metrics. Combining Eq. (18), Eq. (19), and Eq. (21), we have:

$$\lambda^* = \frac{\varepsilon}{2\Delta^*} \leq \frac{\varepsilon}{2 \cdot \sqrt{\phi_{\max}} \cdot \Delta} \quad (24)$$

where λ^* is the rate parameter with the upper bound of $2 \cdot \sqrt{\phi_{\max}} \cdot \Delta$. The rate parameter λ^* governs the exponential

decay of the Laplace-type noise distribution; it controls how rapidly the probability mass concentrates near zero.

In summary, Theorem III.4 establishes the localized Laplace distribution on the conformally transformed metric (\mathcal{M}, g^*) . By specifying the radius r and the rate λ^* within an open geodesic ball $B_r^*(m)$, we obtain the density $K_r^*(z | m)$ that assigns an exponential decay in the conformal distance ρ_{g^*} . Note that Theorem IV.2 shows that if we choose $\lambda^* \leq \frac{\varepsilon}{2\Delta_{\text{loc}}^*}$, where Δ^* is the local sensitivity under ρ_{g^*} , it ensures that the mechanism \mathcal{M} achieves ε -DP. Furthermore, Theorem III.3 ensures the existence of the conformal factor σ , therefore $\phi = e^{2\sigma}$ by solving data-dependent elliptic PDEs, guaranteeing that $g^* = e^{2\sigma}g$ is both smooth and positive-definite. Combined with Corollary III.6, which provides a global bi-Lipschitz relationship between ρ_g and ρ_{g^*} , we obtain explicit upper and lower bounds on geodesic distances in the new metric. Thus, our framework rigorously links the elliptic regularity of the Laplace–Beltrami operator (used to solve for σ) with the construction of ϕ , ensuring that ϕ is smooth and positive under the stated compactness and curvature conditions. However, extending these results to noncompact or higher-dimensional manifolds will require careful analysis of radius dependencies, curvature-driven volume growth, and normalization constants. We leave these challenges to be addressed in future work.

V. THEORETICAL ANALYSIS

In this section, we provide the details to obtain the privatized statistical summary under conformal metrics and the general utility analysis for Conformal-DP.

A. Privatized Fréchet Mean under Conformal Metric

We use Fréchet mean (Karcher mean) for our analysis of the Conformal-DP on Riemannian manifolds as it has been well studied in the literature [63], [64]. Assume a smooth and complete Riemannian manifold \mathcal{M} as previously defined, and dataset $D : \{x_1, \dots, x_n\} \subset \mathcal{M}$, then the Fréchet energy function $F(x)$ on \mathcal{M} is $F(x) = \frac{1}{2n} \sum_{i=1}^n \rho_g^2(x, x_i)$, which reduces in the Euclidean setting $(\mathcal{M}, g) = (\mathbb{R}^d, \langle \cdot, \cdot \rangle)$ to the classical least-squares energy

$$F_{\mathbb{R}^d}(x) = \frac{1}{2n} \sum_{i=1}^n \|x - x_i\|_2^2.$$

The Fréchet mean \bar{x} we are looking for is the minimum point of the function: $\bar{x} = \arg \min_{x \in \mathcal{M}} F(x)$. Karcher's theorem presented in [63] shows that \bar{x} exists and is unique when certain curvature conditions are satisfied, and x_i all fall within the same geodesic convex sphere; and the Riemannian gradient of F is 0 at \bar{x} : $\nabla F(\bar{x}) = \sum_{i=1}^n \log_{\bar{x}}(x_i) = 0$, where $\log_x(\cdot)$ denotes the logarithmic map at x (the local inverse function of \exp_x), i.e., $\log_x(x_i) \in T_{x_i}M$ represents the tangent vector at x corresponding to the geodesic from x to x_i .

Fréchet means on (\mathcal{M}, g^*) . Now, the challenge is to represent Fréchet mean on the conformal metrics. Here, we have the natural generalization of the definition of the mean

TABLE II: A Walk-Through of our proposed Conformal Differential Privacy Algorithm.

Theoretical Basis	Algorithm Design	Notes
Section III-A, f_{data}	$f_{\text{data}}(x) = \mathbf{KDE}(D; x)$.	We use Kernel Density Estimation (KDE) as a discrete $H(x)$ for $\forall x_i \in D$ as f_{data} .
Section III-B, Theorem III.3	$\mathbf{h}_i = -(f_{\text{data}}(x_i) - c)$, in which $c = \frac{1}{N} \sum_{i=1}^N f_{\text{data}}(x_i)$	c is an approximation to $c = \frac{1}{\text{Vol}(M)} \int_M f_{\text{data}} d\mu_g$.
Section III-A, Lemma III.1	$\mathcal{L} \leftarrow \mathbf{D} - \mathbf{W}$, in which: $\mathbf{W}_{i,j} = \begin{cases} \exp(-\ Y_i - Y_j\ _F^2 / \tau^2), & x_j \in \mathbf{KNN}(k; x_i), \\ 0, & \text{otherwise;} \end{cases}$ $\mathbf{D} = \text{diag}\left(\sum_{j \in \mathbf{KNN}(i)} \mathbf{W}_{i,j}\right)$.	\mathcal{L} is an approximation of the Laplace-Beltrami operator Δ_g . \mathbf{D} is the degree matrix and \mathbf{W} is the weighted matrix for harmonic matrix [62]. $Y_i = \log X_i$. We only maintain the relations with the k nearest neighbors for $\forall x_i$, τ is the kernel bandwidth.
$\Delta_g \sigma(x) = H(x)$	Solve $\mathcal{L}\sigma = \mathbf{h}$	Solve the Poisson equation, σ is a vector of N dimensions (conformal scaling vector).
Section III-A	$\phi_i \leftarrow \exp(2\sigma_i)$	Solve conformal factor ϕ .
Section IV-C	$\Delta^* \leftarrow \sqrt{\phi_{\max}} \cdot \Delta$	Calculate the worst-case (local) sensitivity under conformal metric g^* .
Section IV-C, (24)	$\lambda^* \leftarrow \min\left\{\frac{\epsilon}{2\Delta_{\text{loc}}^*(D)}, \lambda_{\max}\right\}$	Calculate the noise rate parameter.
Section IV-B, Theorem IV.2	$\eta_{\text{p}}(D) \sim \mathbb{P}^*(z \mid \eta_{\text{np}}(D), \lambda^*)$ $\propto \exp[-\lambda^* \rho_{g^*}(\eta_{\text{np}}(D), z)]$	Laplacian distribution under ρ^* . We sample $\eta_{\text{p}}(D)$ with Markov Chain Monte Carlo (MCMC) sampling, similar to [18] and [19].

based on [65], we put the original data points $\{x_i\}$ in the conformal metric space (\mathcal{M}, g^*) , $g^* = e^{2\sigma(x)}g$, considering:

$$F^*(x) = \frac{1}{2n} \sum_{i=1}^n \rho_{g^*}^2(x, x_i), \quad (25)$$

we define the new Fréchet means $\bar{x}^* = \arg \min_{m \in \mathcal{M}} F^*(x)$; then we have the logarithmic map:

$$\nabla F^*(\bar{x}^*) = \sum_{i=1}^n \log_{\bar{x}^*}^{(g^*)}(x_i) = 0. \quad (26)$$

However, note that due to the different metrics, the geodesic and exponential mappings are changed accordingly, refer to (11).

B. Utility Analysis

In Section IV, we proposed our data-density-aware *Conformal-DP* mechanism, which follows the strict ϵ -DP form with variable local sensitivity Δ^* based on data density and conformal factor $\phi(x)$. When λ^* is chosen, the distribution of the output of the mechanism $\mathcal{A}(D)$ depends only on the distance $\rho_{g^*}(\eta(D), z)$ to the true mean $f(D)$. That is, the mechanism is isotropic to the g^* metric. To analyze the error, we focus on the distance between the output points and the true Fréchet means. Let us denote a random variable $R^* = \rho_{g^*}(\eta(D), \mathcal{A}(D))$ that represents the conformal distance of the privatized output from the true Fréchet mean f_{np} . Since the density of the distribution is proportional to $\exp(-\lambda^* R^*)$ and isotropic under g^* , we can deduce that the radial probability density function of R^* is similar to the Laplace distribution in Euclidean space. Based on this, we propose Theorem V.1 that gives the explicit form of the Mean Square Error (MSE) of the distance between

the output $\mathcal{A}(D)$ under the g^* metric and the true $\eta(D)$, as follows:

Theorem V.1. *Let the mechanism \mathcal{A} be a Laplace-type perturbation with a linear distance penalty under the conformal metric g^* with rate parameter λ^* , where the output density satisfies: $\mathbb{P}^*(z \mid \eta(D)) \propto \exp(-\lambda^* \rho_{g^*}(\eta(D), z))$. $B_r(z)$ is the geodesic ball that contains the dataset D with the radius of $0 < r < \infty$, $z \in B_r(\mathcal{M})$ is the center of the ball, $\eta(D)$ is the Fréchet mean of D . Under the original metric g , the expected squared distance error between the mechanism output and the Fréchet mean admits the following upper bound:*

$$\mathbb{E}[\rho_g^2(\eta(D), \mathcal{A}(D))] < \frac{16d(d+1)r^2}{\epsilon^2 \eta^2} \frac{\phi_{\max}}{\phi_{\min}}. \quad (27)$$

where d is the dimension of $B_r(\mathcal{M})$, ϵ is the privacy loss budget. The contribution of the extreme points of the integral is infinitesimal, leading to expectations that are strictly less than the upper bound of the theory.

Proof. We present a detailed proof in Appendix G. \square

Theorem V.1 provides an upper-bound for our proposed Conformal-DP mechanism where the worst-case utility scales with $\mathcal{O}(k^4)$, which is consistent with existing results about ϵ -DP on Riemannian manifolds. However, by incorporating data density into the calculation for privacy requirements for different subsets of the dataset (reflected by the conformal factor $\phi_{\min} \leq \phi \leq \phi_{\max}$), we can show an improvement in utility over the Riemannian-Laplace DP [18] while still satisfying ϵ -DP, and better privacy-utility trade-offs than the Tangent-Gaussian DP [19]. We further show in Section VII empirically that the utility of our Conformal-DP mechanism is consistently below the upper bound.

VI. ALGORITHM DESIGN

As discussed in Section III-A (footnote 2), our design of the *Conformal-DP* relies on solving $\Delta_g \sigma(x) = H(x)$, where Δ_g and $H(x)$ are continuous functions over the entire manifold \mathcal{M} . However, there does not always exist an explicit form for the solution to the Laplace–Beltrami operator Δ_g . Therefore, we need to provide discrete approximations for Δ_g and $H(x)$ over $D \in \mathcal{M}^n$ to solve $\sigma(x)$ *discretely* when implementing our proposed *Conformal-DP* mechanism. We derive a solvable Poisson equation for σ as $\mathcal{L}\sigma = \mathbf{h}$. We provide a detailed description of our implementation in Table II.

Approximation for Δ_g . We use a graph Laplacian approximation for Δ_g following Giné and Koltchinskii [66], in which we define the Laplacian matrix \mathcal{L} as $\mathcal{L} = \mathbf{D} - \mathbf{W}$. We define a distance matrix over the original metric $\rho_g(x, y)$, and define the degree matrix \mathbf{D} and the weighted matrix \mathbf{W} on the k -Nearest Neighbors of $\forall x_i \in D$.

Approximation for $H(x)$. Let $h_i = -(f_{\text{data}}(x_i) - c)$, in which $f_{\text{data}}(x)$ is defined as the Kernel Density estimation for x . We use the mean kernel density as an approximation for $c = \frac{1}{\text{Vol}(\mathcal{M})} \int_{\mathcal{M}} f_{\text{data}} d\mu_g$.

VII. EXPERIMENTAL RESULTS

In this section, we demonstrate our experimental results, including the comparison of three different DP mechanisms with synthetic datasets and real-world datasets.

A. Results on Synthetic Datasets

Datasets. Let $\mathcal{M}_{\text{syn}} = \mathcal{B}_1(\mathbf{0}^d)$ be a d -dimensional unit ball ($(d - 1)$ -dimensional hyperplane), to illustrate how our proposed mechanism adjusts privacy perturbation based on data density, we experiment on an heterogeneously distributed data points $x_i \in \mathcal{M}_{\text{syn}}$. We randomly generate x_i as follows: the process starts with a background intensity of μ and generates multiple samples. Existing samples (“parents”) then each generates new samples (“children”) near itself following a Gaussian distribution, the number of children follows a Poisson distribution with an expectation of α . We draw new parent samples on the hyperplane uniformly and generate its children until we reach a total of N samples.

Mechanism. We compare our mechanism to two existing works on privatizing data on the manifolds: Riemannian-Laplace DP proposed by Reimherr et al. in [18] and Tangent-Gaussian DP proposed by Utpala et al. in [19]. The implementation details and hyperparameters are shown in Appendix H.

Metrics. We measure the utility of different DP mechanisms, calculated as:

$$\text{Utility} = \frac{1}{1 + \frac{\rho_g(\eta_p(D), \eta_{np}(D))}{\rho_{g_{\max}}}}$$

In which $\rho_g(\eta(D), \eta_p(D))$ denotes the distance between the privatized (η_p) and non-privatized (η_{np}) Fréchet mean, $\rho_{g_{\max}} = \max[\rho_g(\eta(D), \eta_p(D))]$. A higher utility shows more favorable privacy-utility trade-offs under the same privacy budget ϵ .

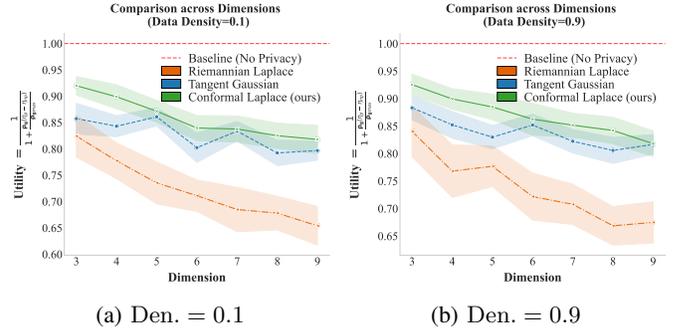


Fig. 2: Comparison of DP utility under different data dimensions ($Dim.$), with density parameter (a) Density = 0.1 and (b) Density = 0.9.

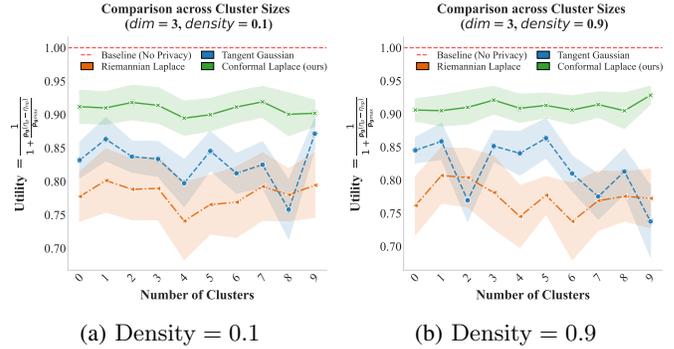


Fig. 3: Comparison of DP utility under different number of clusters with $N = 1000$ samples, with (a) Density = 0.1 and (b) Density = 0.9.

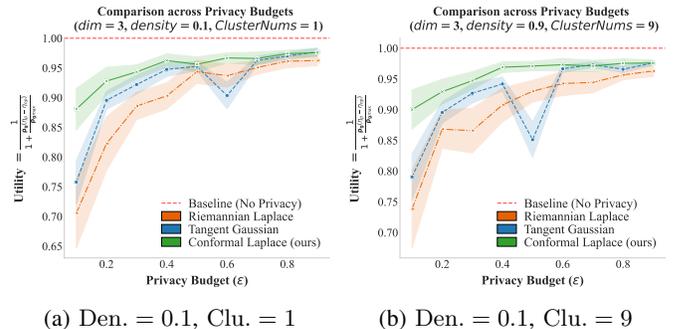


Fig. 4: Comparison of DP utility under different privacy budgets (ϵ), with (a) Dimension = 3, Density = 0.1, Number of Clusters = 1 and (b) Dimension = 3, Density = 0.1, Number of Clusters = 9.

Results. For the synthetic hyperplane, we compare our *Conformal-DP* to Riemannian-Laplace DP [18] and Tangent-Gaussian DP [19] under (a) different dimensions ($3 \leq d \leq 9$); (b) different data density for each cluster, calculated with their standard deviations $0 < std \leq 1$. Note that the small data density parameter is the more dense datasets; and (c) privacy budget $0.1 \leq \epsilon \leq 0.9$. The results are shown in Fig. 2 through 4.

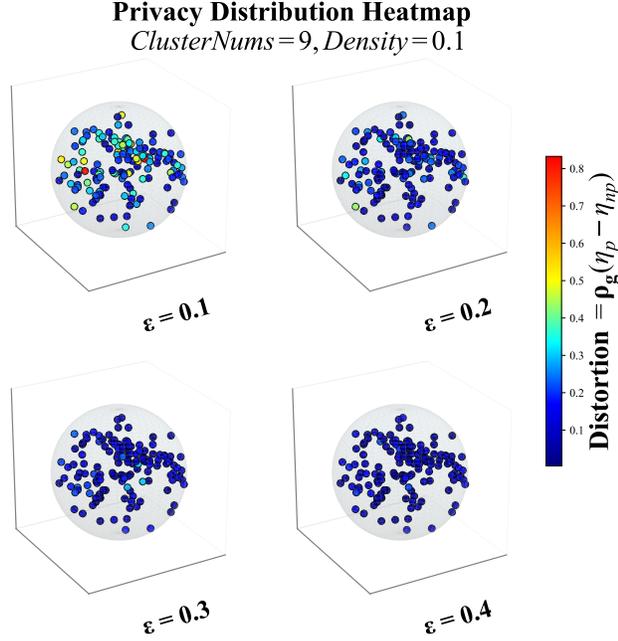


Fig. 5: Heatmap for privacy perturbation distribution on a 3-dimensional hyperplane under Density = 0.1, Number of Clusters = 9.

As demonstrated in our results, our *Conformal*-DP shows superior utility across various scenarios. The dimensionality of the data are crucial to the performance of DP mechanisms. Higher-dimensional data are inherently more sparse and require more perturbations to ensure differential privacy, thus reducing the utility. *Conformal*-DP shows slower and stable utility decrease compared to both the ϵ -DP of [18] and (ϵ, δ) -DP of [19] under different data densities on the hyperplane (Fig. 2a and 2b).

The number of clusters present in $N = 1000$ samples under different density and sizes for each clusters which demonstrate the dynamic nature of the *Conformal*-DP mechanism: while the utilities of Riemannian Laplacian-DP [18] and Tangent Gaussian-DP [19] does not change with the distribution of data samples due to their inherent mechanisms (where utility is only related to the value of ϵ and δ), our proposed *Conformal*-DP is able to balance the data densities of the samples and maintain superior utilities under different number of clusters (Fig. 3a and 3b). Formation of more clusters (which increases local density) does not influence the utility of *Conformal*-DP, which is indicative that our proposed mechanism is able to utilize the heterogeneous distribution and achieve consistent utility under different data densities.

We also demonstrate the utility variations under different privacy loss budgets (Fig. 4a and 4b). As (ϵ) increases, less perturbation is needed, yielding higher utility. Our proposed *Conformal*-DP mechanism consistently achieves superior utility compared to Riemannian Laplacian-DP [18] and Tangent Gaussian-DP [19], which attain similar utility only at larger ϵ

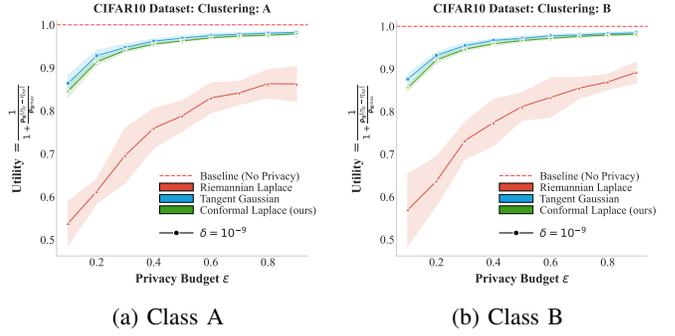


Fig. 6: Utilities of private Fréchet means under varying privacy budgets (ϵ) for classwise CIFAR-10 samples with homogeneous distributions.

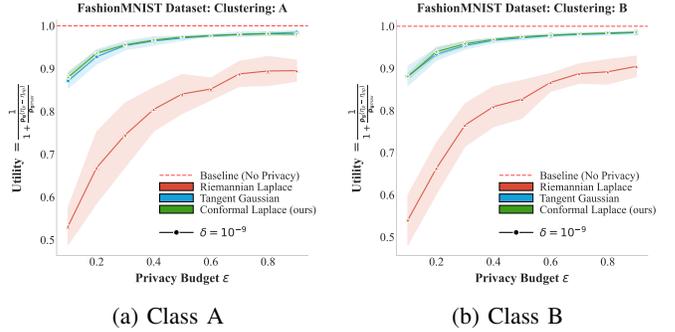


Fig. 7: Utilities of private Fréchet means under varying privacy budgets (ϵ) for classwise Fashion-MNIST samples with homogeneous distributions.

values. Meanwhile, to illustrate the functioning of *Conformal*-DP, we also present heatmaps of privacy distribution at varying privacy budgets (ϵ) in **Figure 5**. Each heatmap visualizes the distance between privatized and non-privatized Fréchet means at local cluster centers, mapped onto a three-dimensional hyperplane. Results highlight that conformal distortions vary heterogeneously according to local data density, with distortions diminishing as the privacy loss budgets increase.

B. Results on Real-world Datasets

Datasets. To emphasize the practical use of the *Conformal*-DP mechanism, we further evaluate the utility of our mechanism for two real-world image datasets with uniform data distribution: **CIFAR-10** [67] and **Fashion-MNIST** [68]. We seek to privatize the Fréchet mean of the data as a descriptor of the dataset $\eta : \mathcal{M}^n \rightarrow \mathcal{M}$. We aim to show that our proposed mechanism is also capable of privatizing image data in Euclidean space for both balanced and imbalanced sampling.

Processing. For image data, we first map them on a Riemannian manifold following Utpala et al. [19], specifically, for an image $\mathcal{I} \in \mathbb{R}^{h \times w \times c}$ represented as $h \times w$ pixels with c channels, we convert it to their covariance descriptor in the form of a $k \times k$ Symmetric Positive Definite (SPD) matrix as:

$$R_\iota(\mathcal{I}) = \frac{1}{|S|} \sum_{\mathbf{x} \in S} (v(\mathcal{I})(\mathbf{x}) - \bar{v})(v(\mathcal{I})(\mathbf{x}) - \bar{v})^T + \iota I. \quad (28)$$

in which $v(\mathcal{I})(\mathbf{x})$ is a pixel-level feature extractor for each pixel $\mathbf{x} = (x, y)$. \bar{v} is the mean of $v(\mathcal{I})(\mathbf{x})$ for $\forall \mathbf{x} \in S$, ι is a small constant. The same as Utpala et al. [19], we define $v(\mathcal{I})(\mathbf{x})$ as:

$$v(\mathcal{I})(\mathbf{x}) = \begin{bmatrix} x, y, \mathcal{I}, |\mathcal{I}_x|, |\mathcal{I}_y|, |\mathcal{I}_{xx}|, |\mathcal{I}_{yy}|, \\ \sqrt{|\mathcal{I}_x|^2 + |\mathcal{I}_y|^2}, \arctan \frac{|\mathcal{I}_x|}{|\mathcal{I}_y|} \end{bmatrix}. \quad (29)$$

Results. We also compare the three DP algorithms – our proposed *Conformal*-DP, Riemannian-Laplace DP [18], and tangent Gaussian DP [19] – on CIFAR-10 and Fashion-MNIST. We show the results for comparisons between utilities of different DP mechanisms on both datasets in Fig. 6 and Fig. 7. In each experiment, we calculate the Fréchet mean $\eta(D)$ and privatized Fréchet mean $\eta_p(D)$ for samples from each class.

In Fig. 6, our proposed *Conformal*-DP mechanism provides a surprising trade-off between privacy guarantee and function utility. The utility between classes differs slightly due to inherent properties of the data in each class (Fig. 6a-6b), but the results over different classes show a consistent trend. Even on *homogeneous* real-world datasets, where our density-aware conformal mechanism was *not* primarily intended to shine, we still observe a marked advantage over the Riemannian-Laplace mechanism [18]. In particular, Theorem V.1 establishes a tighter theoretical upper-bound on the excess risk, which matches our empirical findings. More strikingly, our method attains similar utilities as the tangent-Gaussian mechanism [19], yet without relying on the (ϵ, δ) relaxation, we retain a strict ϵ -DP guarantee. These results demonstrate that the proposed Conformal-DP is broadly applicable: it adapts seamlessly to both heterogeneous and homogeneous datasets, consistently delivering high utility under strict privacy guarantees.

VIII. RELATED WORK

Manifold DP. Manifold data has been widely studied in statistical settings. Earlier works by Fletcher et al. [69], Pennec et al. [65], and Dryden et al. [3] have introduced the importance and techniques for statistical methods on manifold data. Particularly, Fisher et al. [70] have emphasized the benefit of a spherical manifold’s representation of data. Our idea could potentially utilize the manifold intrinsic properties, which provide better alignment with varying data distributions.

More recently, various works have extended Laplace, Gaussian, and variants of classical DP mechanisms to general manifolds [18], [29], [37], [71]. Reimherr et al. first extended the Laplace / K-norm DP mechanism to general manifolds by leveraging the intrinsic geometric properties of the manifold in the privacy perturbation. Similarly, Soto et al. [36] showed

improved utility by utilizing the k -norm gradient mechanism of Laplace DP on Riemannian manifolds. We compare our proposed Conformal-DP mechanism as an extension of Laplacian DP on Riemannian manifolds to show how DP can incorporate data density in privacy perturbation on manifolds while preserving the ϵ -DP guarantee.

IX. CONCLUSION

In this paper, we propose *Conformal*-DP, a novel data-density-aware differential privacy mechanism on Riemannian manifolds using conformal transformations. By introducing conformal metrics, our method adaptively adjusts privacy perturbations based on local data density. Theoretical and empirical analyses demonstrate that *Conformal*-DP achieves an improved privacy-utility trade-off for both homogeneous and particularly heterogeneous data distributions when privatizing Fréchet means on Riemannian manifolds. A concrete application of *Conformal*-DP is in diffusion-tensor magnetic resonance imaging (DT-MRI), where spatially varying tissue microstructure densities yield highly heterogeneous data: by adapting the noise scale to local kernel-density estimates, *Conformal*-DP injects minimal perturbation in regions of high sampling density—thereby preserving critical diagnostic metrics such as fractional anisotropy or mean diffusivity, while still guaranteeing ϵ -differential privacy for each voxel’s diffusion measurement.

Future research includes extending this method to non-compact manifolds (open spaces) or manifolds with boundaries, requiring careful management of volume alterations and boundary conditions. Another significant direction involves scaling the approach to higher-dimensional manifolds. Additionally, extending privacy protections to more complex statistical measures beyond basic statistics remains an important open challenge.

ACKNOWLEDGMENTS

This research is supported in part by the University of Pittsburgh Center for Research Computing and Data, RRID: SCR_022735, through the resources provided. Specifically, this work used the HTC cluster, which is supported by NIH award number S10OD028483. Rahimian is partially supported by the National Science Foundation under Grant No. 2318844.

REFERENCES

- [1] X. Pennec, S. Sommer, and T. Fletcher, *Riemannian Geometric Statistics in Medical Image Analysis*. Academic Press, 2019.
- [2] I. L. Dryden, *Statistical Analysis on High-Dimensional Spheres and Shape Spaces*. Chichester, UK: John Wiley & Sons, 2005.
- [3] I. L. Dryden, A. Koloydenko, and D. Zhou, “Non-euclidean statistics for covariance matrices, with applications to diffusion tensor imaging,” *The Annals of Applied Statistics*, vol. 3, no. 3, pp. 1102–1123, 2009.
- [4] M. Belkin, P. Niyogi, and V. Sindhwani, “Manifold regularization: A geometric framework for learning from labeled and unlabeled examples,” *The Journal of Machine Learning Research*, vol. 7, pp. 2399–2434, 2006.
- [5] P. Niyogi, “Manifold regularization and semi-supervised learning: Some theoretical analyses,” *The Journal of Machine Learning Research*, vol. 14, no. 1, pp. 1229–1250, 2013.

- [6] G. Cheng and B. C. Vemuri, "A novel dynamic system in the space of spd matrices with applications to appearance tracking," *SIAM Journal on Imaging Sciences*, vol. 6, no. 1, pp. 592–615, 2013.
- [7] P. Turaga, A. Veeraraghavan, and R. Chellappa, "Statistical analysis on stiefel and grassmann manifolds with applications in computer vision," in *2008 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2008, pp. 1–8.
- [8] P. K. Turaga and A. Srivastava, *Riemannian Computing in Computer Vision*. Springer, 2016, vol. 1.
- [9] C. Dwork, "Differential privacy," in *ICALP 2006: Automata, Languages and Programming, Part II*, ser. Lecture Notes in Computer Science, vol. 4052. Berlin, Heidelberg: Springer, 2006, pp. 1–12.
- [10] C. Dwork, G. N. Rothblum, and S. Vadhan, "Boosting and differential privacy," in *2010 IEEE 51st Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE. IEEE, 2010, pp. 51–60.
- [11] C. Dwork, "Differential privacy: A survey of results," in *Theory and Applications of Models of Computation*, M. Agrawal, D. Du, Z. Duan, and A. Li, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 1–19.
- [12] I. Mironov, "Rényi differential privacy," in *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*. IEEE, 2017, pp. 263–275.
- [13] C. Dwork and G. N. Rothblum, "Concentrated differential privacy," 2016, arXiv:1603.01887.
- [14] R. Hall, A. Rinaldo, and L. Wasserman, "Random differential privacy," 2011, arXiv:1112.2680.
- [15] J. B. Tenenbaum, V. d. Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, 2000.
- [16] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, no. 5500, pp. 2323–2326, 2000.
- [17] J. Xie, G. Dai, F. Zhu, E. K. Wong, and Y. Fang, "Deepshape: Deep-learned shape descriptor for 3d shape retrieval," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 7, pp. 1335–1345, 2016.
- [18] M. Reimherr, K. Bharath, and C. Soto, "Differential privacy over riemannian manifolds," *Advances in Neural Information Processing Systems*, vol. 34, pp. 12 292–12 303, 2021.
- [19] S. Utpala, P. Vepakomma, and N. Miolane, "Differentially private fréchet mean on the manifold of symmetric positive definite (spd) matrices with log-euclidean metric," 2022, arXiv:2208.04245.
- [20] J. Zhao, J. Zhang, and H. V. Poor, "Dependent differential privacy for correlated data," in *2017 IEEE Globecom Workshops*. IEEE, 2017, pp. 1–7.
- [21] Y. Zhao, J. T. Du, and J. Chen, "Scenario-based adaptations of differential privacy: A technical survey," *ACM Computing Surveys*, vol. 56, no. 8, pp. 1–39, 2024.
- [22] L. A. Dunning and R. Kresman, "Privacy preserving data sharing with anonymous id assignment," *IEEE Transactions on Information Forensics and Security*, vol. 8, no. 2, pp. 402–413, 2012.
- [23] Y. Liang and K. Yi, "Smooth sensitivity for geo-privacy," in *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 333–347.
- [24] M. Obata, "Conformal transformations of riemannian manifolds," *Journal of Differential Geometry*, vol. 4, no. 3, pp. 311–333, 1970.
- [25] R. Schoen, "Conformal deformation of a riemannian metric to constant scalar curvature," *Journal of Differential Geometry*, vol. 20, no. 2, pp. 479–495, 1984.
- [26] P. Aviles and R. C. McOwen, "Conformal deformation to constant negative scalar curvature on noncompact riemannian manifolds," *Journal of Differential Geometry*, vol. 27, no. 2, pp. 225–239, 1988.
- [27] P. Francesco, P. Mathieu, and D. Sénéchal, *Conformal Field Theory*. Springer Science & Business Media, 2012.
- [28] C. Dwork, A. Roth *et al.*, "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.
- [29] Y. Jiang, X. Chang, Y. Liu, L. Ding, L. Kong, and B. Jiang, "Gaussian differential privacy on riemannian manifolds," *Advances in Neural Information Processing Systems*, vol. 36, pp. 14 665–14 684, 2023.
- [30] P. Ginsparg, "Applied conformal field theory," arXiv preprint hep-th/9108028, 1990.
- [31] I. Katsman, E. Chen, S. Holalkere, A. Asch, A. Lou, S. N. Lim, and C. M. De Sa, "Riemannian residual neural networks," in *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds., vol. 36. Curran Associates, Inc., 2023, pp. 63 502–63 514.
- [32] R. E. Greene, "Review of "riemannian geometry" by s. gallot, d. hulin, and j. lafontaine," 1989, book review.
- [33] S. Lang, *Introduction to Differentiable Manifolds*. Springer Science & Business Media, 2006.
- [34] A. Srivastava and E. P. Klassen, *Functional and Shape Data Analysis*. Springer, 2016, vol. 1.
- [35] J. M. Lee, *Introduction to Riemannian Manifolds*. Springer, 2018, vol. 2.
- [36] C. Soto, K. Bharath, M. Reimherr, and A. Slavković, "Shape and structure preserving differential privacy," *Advances in Neural Information Processing Systems*, vol. 35, pp. 24 693–24 705, 2022.
- [37] J. Dong, A. Roth, and W. J. Su, "Gaussian differential privacy," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 84, no. 1, pp. 3–37, 2022.
- [38] T. del Castillo and Tominich, *Differentiable Manifolds*. Springer, 2020.
- [39] A. Ottazzi, "The liouville theorem for conformal maps: old and new," Seminario Dottorato 2010/11, University of Padua, 2010, available via Citeseer.
- [40] S. Weinberg, *Gravitation and Cosmology: Principles and Applications of the General Theory of Relativity*. John Wiley & Sons, 2013.
- [41] L. Wasserman and S. Zhou, "A statistical framework for differential privacy," *Journal of the American Statistical Association*, vol. 105, no. 489, pp. 375–389, 2010.
- [42] J. Awan, A. Kenney, M. Reimherr, and A. Slavković, "Benefits and pitfalls of the exponential mechanism with applications to Hilbert spaces and functional PCA," in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 09–15 Jun 2019, pp. 374–384.
- [43] E. Parzen, "On estimation of a probability density function and mode," *The Annals of Mathematical Statistics*, vol. 33, no. 3, pp. 1065–1076, 1962. [Online]. Available: <http://www.jstor.org/stable/2237880>
- [44] H. Urakawa, "Geometry of laplace-beltrami operator on a complete riemannian manifold," *Progress in Differential Geometry*, vol. 22, pp. 347–406, 1993.
- [45] Z. Li, Z. Shi, and J. Sun, "Point integral method for solving poisson-type equations on manifolds from point clouds with convergence guarantees," *Communications in Computational Physics*, vol. 22, no. 1, pp. 228–258, 2017.
- [46] L. Beck, *Elliptic Regularity Theory*, ser. Lecture Notes of the Unione Matematica Italiana. Cham: Springer, 2016, vol. 19.
- [47] X. Fernández-Real and X. Ros-Oton, *Regularity Theory for Elliptic PDE*. EMS Press, 2023.
- [48] D. Gilbarg and N. S. Trudinger, *Elliptic Partial Differential Equations of Second Order*, ser. Classics in Mathematics. Berlin, Heidelberg: Springer Berlin Heidelberg, jan 1997.
- [49] J. Jost and J. Jost, *Riemannian Geometry and Geometric Analysis*. Springer, 2008, vol. 42005.
- [50] G. Schwarz, *Hodge Decomposition-A Method for Solving Boundary Value Problems*. Springer, 2006.
- [51] E. Calabi, "An extension of e. Hopf's maximum principle with an application to riemannian geometry," *Duke Mathematical Journal*, vol. 25, no. 1, pp. 45–56, March 1958.
- [52] A. Mirshani, M. Reimherr, and A. Slavković, "Formal privacy for functional data with Gaussian perturbations," in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 09–15 Jun 2019, pp. 4595–4604.
- [53] R. Grone, R. Merris, and V. S. Sunder, "The laplacian spectrum of a graph," *SIAM Journal on Matrix Analysis and Applications*, vol. 11, no. 2, pp. 218–238, 1990.
- [54] B. Gnedenko, "Sur la distribution limite du terme maximum d'une serie aleatoire," *Annals of Mathematics*, vol. 44, no. 3, pp. 423–453, 1943.
- [55] S. Banach and H. Steinhaus, "Sur le principe de la moyenne et la limite de la moyenne d'une fonction," *Studia Mathematica*, vol. 2, pp. 50–68, 1930.
- [56] J. Matoušek, "Bi-lipschitz embeddings into low-dimensional euclidean spaces," *Commentationes Mathematicae Universitatis Carolinae*, vol. 31, no. 3, pp. 589–600, 1990.

- [57] I. Chavel, *Riemannian Geometry: A Modern Introduction*, ser. Cambridge Tracts in Mathematics. Cambridge: Cambridge University Press, 1995, vol. 108.
- [58] P. Topping, *Lectures on the Ricci Flow*, ser. London Mathematical Society Lecture Note Series. Cambridge: Cambridge University Press, 2006, vol. 325.
- [59] S. Rosenberg, *The Laplacian on a Riemannian Manifold: An Introduction to Analysis on Manifolds*, ser. Cambridge Tracts in Mathematics. Cambridge: Cambridge University Press, 1997, vol. 31.
- [60] B. Balle and Y.-X. Wang, “Improving the Gaussian mechanism for differential privacy: Analytical calibration and optimal denoising,” in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, J. Dy and A. Krause, Eds., vol. 80. PMLR, 10–15 Jul 2018, pp. 394–403.
- [61] P. Petersen, *Riemannian Geometry*. Springer, 2006, vol. 171.
- [62] M. Reuter, S. Biasotti, D. Giorgi, G. Patanè, and M. Spagnuolo, “Discrete laplace–beltrami operators for shape analysis and segmentation,” *Computers & Graphics*, vol. 33, no. 3, pp. 381–390, 2009.
- [63] H. Karcher, “Riemannian center of mass and mollifier smoothing,” *Communications on pure and applied mathematics*, vol. 30, no. 5, pp. 509–541, 1977.
- [64] W. S. Kendall, “Probability, convexity, and harmonic maps with small image i: uniqueness and fine existence,” *Proceedings of the London Mathematical Society*, vol. 3, no. 2, pp. 371–406, 1990.
- [65] X. Pennec, “Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements,” *Journal of Mathematical Imaging and Vision*, vol. 25, pp. 127–154, 2006.
- [66] E. Giné and V. Koltchinskii, “Empirical graph laplacian approximation of laplace–beltrami operators: Large sample results,” in *High Dimensional Probability*, ser. Lecture Notes–Monograph Series. Institute of Mathematical Statistics, 2006, vol. 49, pp. 238–259.
- [67] A. Krizhevsky, “Learning multiple layers of features from tiny images,” University of Toronto, Technical Report TR-2009, 2009.
- [68] H. Xiao, K. Rasul, and R. Vollgraf, “Fashion-mnist: A novel image dataset for benchmarking machine learning algorithms,” 2017, arXiv:1708.07747.
- [69] P. T. Fletcher, C. Lu, S. M. Pizer, and S. Joshi, “Principal geodesic analysis for the study of nonlinear statistics of shape,” *IEEE Transactions on Medical Imaging*, vol. 23, no. 8, pp. 995–1005, 2004.
- [70] N. I. Fisher, *Statistical Analysis of Circular Data*. Cambridge University Press, 1995.
- [71] M. Bun and T. Steinke, “Concentrated differential privacy: Simplifications, extensions, and lower bounds,” in *Proceedings, Part I, of the 14th International Conference on Theory of Cryptography - Volume 9985*. Berlin, Heidelberg: Springer-Verlag, 2016, p. 635–658.
- [72] X.-J. Wang *et al.*, “Schauder estimates for elliptic and parabolic equations,” *Chinese Annals of Mathematics-Series B*, vol. 27, no. 6, p. 637, 2006.
- [73] L. C. Evans, *Partial Differential Equations*. American Mathematical Society, 2022, vol. 19.
- [74] P. Pucci and J. B. Serrin, *The Maximum Principle*. Springer Science & Business Media, 2007, vol. 73.
- [75] E. W. Stacy, “A generalization of the gamma distribution,” *The Annals of Mathematical Statistics*, vol. 33, no. 3, pp. 1187–1192, 1962. [Online]. Available: <http://www.jstor.org/stable/2237889>
- [76] N. Miolane, A. L. Brigant, J. Mathe, B. Hou, N. Guigui, Y. Thanwerdas, S. Heyder, O. Peltre, N. Koep, H. Zaatiti, H. Hajri, Y. Cabanes, T. Gerald, P. Chauchat, C. Shewmake, B. Kainz, C. Donnat, S. P. Holmes, and X. Pennec, “Geomstats: A python package for riemannian geometry in machine learning,” 2020, arXiv:2004.04667.
- [77] N. Metropolis and S. Ulam, “The monte carlo method,” *Journal of the American Statistical Association*, vol. 44, no. 247, pp. 335–341, 1949.

APPENDIX

A. Proof for Theorem III.3

Proof. Note, we consider a typical second-order strongly elliptic operator Ω_g in local coordinates $\{x_1, \dots, x_n\}$ on the

Riemannian manifold (\mathcal{M}, g) which is in divergence form:

$$\begin{aligned} \Omega_g[\sigma](x) &= \frac{1}{\sqrt{\det(g)}} \frac{\partial}{\partial x^i} \left(\sqrt{\det(g)} A^{ij}(x, \sigma(x)) \frac{\partial \sigma}{\partial x^j}(x) \right) + \\ &\quad B^i(x, \sigma(x)) \frac{\partial \sigma}{\partial x^i}(x) + C(x, \sigma(x)) \sigma(x) \\ &= \frac{1}{\sqrt{\det(g)}} \frac{\partial}{\partial x^i} \left(\sqrt{\det(g)} g^{ij} \frac{\partial \sigma}{\partial x^j} \right) \\ &= \Delta_g \sigma(x) \end{aligned}$$

where $A^{ij}(x, \mu)$ is the leading-order coefficient with the property of positive-definite; in our case, $A^{ij}(x, \mu) = g^{ij}$; $B^i(x, \mu)$ and $C(x, \mu)$ are lower-order terms, which can also depend on μ and its derivatives if the operator is nonlinear, in our case, $B^i = C = 0$. Since Ω_g is strongly elliptic and its coefficients (as well as H) are smooth, using Schauder estimates [72] guarantees that any weak solution is in $C^\infty(\mathcal{M})$, thus σ is in fact smooth. In our application, $\Omega_g = \Delta_g$ (a special case of such elliptic operators), so a solution σ to $\Delta_g \sigma = H$ exists and is C^∞ by elliptic regularity [48]. Note that in the theorem, we assume $\phi(x) = e^{2\sigma(x)}$, which ensures $\phi(x) > 0$ for every $x \in \mathcal{M}$. Since $\sigma \in C^\infty(\mathcal{M})$, it follows that $\phi \in C^\infty(\mathcal{M})$. Because ϕ is strictly positive, each tensor $g(x)$ remains positive-definite on the tangent space $T_x \mathcal{M}$, and the smoothness of ϕ and g implies $g(x)$ is itself in $C^\infty(\mathcal{M})$. This confirms that g^* is a smooth, positive Riemannian metric, thereby proving the theorem’s assertion that the conformal metric remains smooth and positive-definite. \square

B. Proof for Theorem III.4

Proof. Let \mathcal{M} is compact and the Laplacian Δ_g is a classical, self-adjoint second-order elliptic operator, the Poisson-type equation

$$-\Delta_g \sigma = f_{\text{data}}$$

The constraint $\int_{\mathcal{M}} \sigma d\mu_g = 0$ admits a unique (up to additive constants) solution. The additional condition $\int_{\mathcal{M}} \sigma d\mu_g = 0$ fixes the constant part uniquely. Since Δ_g has a discrete spectrum $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \dots$ on a compact manifold, the eigenfunctions $\{\varphi_i\}_{i=0}^\infty$ form an orthonormal basis in $L^2(\mathcal{M})$. We write $-\Delta_g \varphi_i = \lambda_i \varphi_i$, $\int_{\mathcal{M}} \varphi_i \varphi_j d\mu_g = \delta_{ij}$. Here, $\lambda_1 > 0$ is the smallest positive eigenvalue, often called the *spectral gap* [53].

Since f_{data} has zero mean, it is orthogonal to the constant mode φ_0 . Hence, we can expand $f_{\text{data}}(x) = \sum_{i=1}^\infty a_i \varphi_i(x)$, where $a_i = \int_{\mathcal{M}} f_{\text{data}}(x) \varphi_i(x) d\mu_g$. The solution σ then has the corresponding expansion $\sigma(x) = -\sum_{i=1}^\infty \frac{a_i}{\lambda_i} \varphi_i(x)$, where we omit $i = 0$ because of the zero-mean condition on σ .

To prove $\|\sigma\|_{L^\infty(\mathcal{M})} \leq \frac{1}{\lambda_1} \|f_{\text{data}}\|_{L^\infty(\mathcal{M})}$, we invoke the fact that $(-\Delta_g)^{-1}$ acts as a bounded linear operator on the space of mean-zero $C^0(\mathcal{M})$ functions, with operator norm of $1/\lambda_1$. A direct argument [48], [73] shows that if $-\Delta_g u = g$ and $\int_{\mathcal{M}} u = 0$, then $\|u\|_{L^\infty(\mathcal{M})} \leq \frac{1}{\lambda_1} \|g\|_{L^\infty(\mathcal{M})}$. Applying this to $u = \sigma$ and $g = f_{\text{data}}$ immediately yields $\|\sigma\|_{L^\infty(\mathcal{M})} \leq \frac{1}{\lambda_1} \|f_{\text{data}}\|_{L^\infty(\mathcal{M})}$.

Then we define $f_{\min} = \min_{x \in \mathcal{M}} f_{\text{data}}(x)$ and $f_{\max} = \max_{x \in \mathcal{M}} f_{\text{data}}(x)$. Because $\int_{\mathcal{M}} f_{\text{data}} d\mu_g = 0$, one generally expects $f_{\min} \leq 0 \leq f_{\max}$ unless f_{data} is identically zero.

Let $p \in \mathcal{M}$ be a point where $\sigma(p) = \max_{x \in \mathcal{M}} \sigma(x)$. By the strong maximum principle for elliptic equations [74, Chapter 2], at p we must have $-\Delta_g \sigma(p) = f_{\text{data}}(p)$. But p is a maximum of σ , so heuristically $\Delta_g \sigma(p) \leq 0$ (since the Laplacian at a maximum is non-positive). Hence $f_{\text{data}}(p) \leq 0$, which implies $f_{\text{data}}(p) \geq f_{\min}$ but also $f_{\min} \leq 0$, w.r.t, $\Delta_g \sigma(p) = -f_{\text{data}}(p) \geq -f_{\min}$. Applying the known relationship between λ_1 and the Laplacian, we deduce $\sigma_{\max} = \sigma(p) \leq -\frac{f_{\min}}{\lambda_1}$.

Similarly, let $q \in \mathcal{M}$ be the point where $\sigma(q) = \min_{x \in \mathcal{M}} \sigma(x)$. Then a parallel argument via the minimum principle indicates $-\Delta_g \sigma(q) = f_{\text{data}}(q)$, and at a minimum of σ , $\Delta_g \sigma(q) \geq 0$. Hence $f_{\text{data}}(q) \geq 0$, which implies $f_{\text{data}}(q) \leq f_{\max}$ and $f_{\max} \geq 0$. Thus

$$\sigma_{\min} = \sigma(q) \geq -\frac{f_{\max}}{\lambda_1}.$$

Combining them we have:

$$-\frac{f_{\max}}{\lambda_1} \leq \sigma(x) \leq -\frac{f_{\min}}{\lambda_1} \quad \forall x \in \mathcal{M}.$$

The two main results of Theorem III.4 thus follow:

$$\|\sigma\|_{L^\infty(\mathcal{M})} \leq \frac{1}{\lambda_1} \|f_{\text{data}}\|_{L^\infty(\mathcal{M})},$$

$$-\frac{f_{\max}}{\lambda_1} \leq \sigma(x) \leq -\frac{f_{\min}}{\lambda_1} \quad \forall x \in \mathcal{M}.$$

This completes the proof. \square

C. Proof for Corollary III.6

Proof. Since ϕ is smooth on compact M , it attains a finite maximum and positive minimum ϕ , ϕ_{\min} and ϕ_{\max} from Lemma III.5, then it is possible to transfer the inequalities to geodesic distances. For lower bound transfer, $L_{g^*}(\gamma) \geq \sqrt{\phi_{\min}} L_g(\gamma) \Rightarrow \inf_{\gamma} L_{g^*}(\gamma) \geq \sqrt{\phi_{\min}} \inf_{\gamma} L_g(\gamma)$; thus we have:

$$\rho_{g^*}(x, y) \geq \sqrt{\phi_{\min}} \rho_g(x, y).$$

For upper bound transfer, $L_{g^*}(\gamma) \leq \sqrt{\phi_{\max}} L_g(\gamma) \Rightarrow \inf_{\gamma} L_{g^*}(\gamma) \leq \sqrt{\phi_{\max}} \inf_{\gamma} L_g(\gamma)$; thus we have:

$$\rho_{g^*}(x, y) \leq \sqrt{\phi_{\max}} \rho_g(x, y).$$

In particular, distances cannot shrink or stretch by more than factors $\sqrt{\phi_{\min}}$ and $\sqrt{\phi_{\max}}$, respectively. Therefore, $\sqrt{\phi_{\min}} \rho_g(x, y) \leq \rho_{g^*}(x, y) \leq \sqrt{\phi_{\max}} \rho_g(x, y)$; for any case, these estimates hold for the entire manifold and for all paths. This implies that the new distance ρ_{g^*} is bi-Lipschitz equivalent to the original distance ρ_g . This implies that no pair of points gets closer by more than a factor $\sqrt{\phi_{\min}}$ nor farther by more than $\sqrt{\phi_{\max}}$ under the conformal transformation. \square

D. Proof for Theorem IV.2

Proof. Let $\eta \in \mathcal{M}$ be the chosen center point (Fréchet Mean), which is determined by the dataset D . We aim to construct a localized Laplace-type distribution that places exponential-decay noise around η , truncated to a geodesic ball of radius r in the conformal metric g^* . We define an unnormalized Laplace-type kernel: $\tilde{K}_r^*(z | \eta) := \exp[-\lambda^* \rho_{g^*}(\eta, z)] \mathbf{1}_{\{\rho_{g^*}(\eta, z) < r\}}$, $z \in \mathcal{M}$ where $\mathbf{1}_{\{\rho_{g^*}(\eta, z) < r\}}$ is the indicator function that is 1 if z lies inside the conformal geodesic ball $B_r^*(\eta)$ and 0 otherwise; ρ_{g^*} denotes the geodesic distance induced by g^* ; $\lambda^* > 0$ is a rate parameter. By definition, $\tilde{K}_r^*(z | \eta)$ is zero outside $B_r^*(\eta)$ and exponential inside $B_r^*(\eta)$. It is not yet a probability distribution, as its integral need not be 1. We next compute the integral over the entire manifold \mathcal{M} :

$$\begin{aligned} \int_{\mathcal{M}} \tilde{K}_r^*(z | \eta) d\mu_{g^*}(z) &= \int_{B_r^*(\eta)} \exp[-\lambda^* \rho_{g^*}(\eta, z)] d\mu_{g^*}(z) \\ &\equiv C_r(\eta, \lambda^*), \end{aligned}$$

Since $\tilde{K}_r^*(z | \eta) = 0$ outside $B_r^*(\eta)$, effectively

$$C_r(\eta, \lambda^*) = \int_{B_r^*(\eta)} \exp[-\lambda^* \rho_{g^*}(\eta, z)] d\mu_{g^*}(z). \quad (30)$$

Because $\rho_{g^*}(\eta, z)$ is finite on the geodesic ball $B_r^*(\eta)$ and the exponential function is integrable on a complete Riemannian manifold, this integral is finite and strictly positive:

$$0 < C_r(\eta, \lambda^*) < +\infty.$$

To make \tilde{K}_r^* into a proper probability distribution, we normalize it by the positive constant $Z_r(\eta, \lambda^*)$. We thus define:

$$\mathbb{P}_r^*(z | \eta) := \frac{\tilde{K}_r^*(z | \eta)}{C_r(\eta, \lambda^*)} = \frac{\exp[-\lambda^* \rho_{g^*}(\eta, z)] \mathbf{1}_{\{\rho_{g^*}(\eta, z) < r\}}}{\int_{B_r^*(\eta)} \exp[-\lambda^* \rho_{g^*}(\eta, u)] d\mu_{g^*}(u)}$$

thus:

$$\mathbb{P}_r^*(z | \eta) = \begin{cases} \frac{\exp[-\lambda^* \rho_{g^*}(\eta, z)]}{\int_{B_r^*(\eta)} \exp[-\lambda^* \rho_{g^*}(\eta, u)] d\mu_{g^*}(u)}, & \rho_{g^*}(\eta, z) < r, \\ 0, & \rho_{g^*}(\eta, z) \geq r. \end{cases} \quad (31)$$

We check that its integral over all of \mathcal{M} is 1. Note that $\mathbb{P}_r^*(z | \eta)$ is zero outside $B_r^*(\eta)$, so:

$$\begin{aligned} &\int_{\mathcal{M}} \mathbb{P}_r^*(z | \eta) d\mu_{g^*}(z) \\ &= \int_{B_r^*(\eta)} \frac{\exp[-\lambda^* \rho_{g^*}(\eta, z)]}{\int_{B_r^*(\eta)} \exp[-\lambda^* \rho_{g^*}(\eta, u)] d\mu_{g^*}(u)} d\mu_{g^*}(z) \\ &= \frac{\int_{B_r^*(\eta)} \exp[-\lambda^* \rho_{g^*}(\eta, z)] d\mu_{g^*}(z)}{\int_{B_r^*(\eta)} \exp[-\lambda^* \rho_{g^*}(\eta, u)] d\mu_{g^*}(u)} \\ &= \frac{C_r^*(\eta, \lambda^*)}{C_r^*(\eta, \lambda^*)} = 1. \end{aligned}$$

We define the localized conformal Laplace mechanism $\mathcal{A}_{r, \lambda^*}(\eta)$ to be a random variable Z taking values in \mathcal{M} with probability density $\mathbb{P}_r^*(\cdot | \eta)$. Concretely: $\mathbb{P}(Z \in \mathcal{A}) =$

$\int_{\mathcal{A}} K_r^*(z | \eta) d\mu_{g^*}(z)$, $\forall \mathcal{A} \subset \mathcal{M}$ measurable. This completes the construction of a localized (radius = r) Laplace-type noise mechanism under the conformal metric g^* . \square

E. Proof for Theorem IV.4

Proof. By definition of the random mechanism \mathcal{A} , we have:

$$\mathbb{P}^*(z | \eta(D), \lambda^*) = \frac{1}{C(\eta(D), \lambda^*)} \exp[-\lambda^* \rho_{g^*}(\eta(D), z)],$$

and likewise:

$$\mathbb{P}^*(z | \eta(D'), \lambda^*) = \frac{1}{C(\eta(D'), \lambda^*)} \exp[-\lambda^* \rho_{g^*}(\eta(D'), z)].$$

Hence,

$$\frac{\mathbb{P}^*(z | \eta(D), \lambda^*)}{\mathbb{P}^*(z | \eta(D'), \lambda^*)} = \frac{\exp[-\lambda^* \rho_{g^*}(\eta(D), z)]}{\exp[-\lambda^* \rho_{g^*}(\eta(D'), z)]} \times \frac{C(\eta(D'), \lambda^*)}{C(\eta(D), \lambda^*)}.$$

Taking the natural logarithm of both sides directly yields:

$$\begin{aligned} L_{D, D'}(z) &:= \log \frac{\mathbb{P}^*(z | \eta(D), \lambda^*)}{\mathbb{P}^*(z | \eta(D'), \lambda^*)} \\ &= -\lambda^* [\rho_{g^*}(\eta(D), z) - \rho_{g^*}(\eta(D'), z)] \\ &\quad + \log \frac{C(\eta(D'), \lambda^*)}{C(\eta(D), \lambda^*)} \end{aligned}$$

This completes the derivation of $L_{D, D'}(z)$. For simplicity, we denote $x = \eta(D)$, $y = \eta(D')$, $a = \rho_{g^*}(x, z)$, $b = \rho_{g^*}(y, z)$, and u as a dummy variable. For $-\lambda^* [\rho_{g^*}(x, z) - \rho_{g^*}(y, z)]$, we then have $-\lambda^*[a - b] = \lambda^*[b - a]$, based on Triangle Inequality theorem: $\rho_{g^*}(y, z) \leq \rho_{g^*}(y, x) + \rho_{g^*}(x, z) \Rightarrow b \leq \rho_{g^*}(x, y) + a$; thus, we have $b - a \leq \rho_{g^*}(x, y)$, then $\lambda^*[b - a] \leq \lambda^* \rho_{g^*}(x, y) \Rightarrow -\lambda^* [\rho_{g^*}(x, z) - \rho_{g^*}(y, z)] \leq \lambda^* \rho_{g^*}(x, y)$. We then calculate the upper bound of the normalization constant ratio $\log \left[\frac{Z(y)}{Z(x)} \right]$.

$$\begin{aligned} C(y) &= \int_{\mathcal{M}} \exp[-\lambda^* \rho_{g^*}(y, u)] d\mu_{g^*}(u). \\ C(x) &= \int_{\mathcal{M}} \exp[-\lambda^* \rho_{g^*}(x, u)] d\mu_{g^*}(u). \end{aligned}$$

Based on Triangle Inequality, we have $\rho_{g^*}(y, u) \leq \rho_{g^*}(y, x) + \rho_{g^*}(x, u)$, carrying this into the exponent gives

$$\begin{aligned} \exp[-\lambda^* \rho_{g^*}(y, u)] &\leq \exp[-\lambda^* (\rho_{g^*}(x, u) - \rho_{g^*}(y, x))] \\ &= \exp[-\lambda^* \rho_{g^*}(x, u)] \times \exp[+\lambda^* \rho_{g^*}(y, x)], \end{aligned}$$

Next, we have

$$\exp[-\lambda^* \rho_{g^*}(y, u)] \leq \exp[-\lambda^* \rho_{g^*}(x, u)] \times \exp[\lambda^* \rho_{g^*}(y, x)].$$

Bring to $Z(y)$'s integration, we have:

$$\begin{aligned} C(y) &= \int_{\mathcal{M}} \exp[-\lambda^* \rho_{g^*}(y, u)] d\mu_{g^*}(u) \\ &\leq \int_{\mathcal{M}} \exp[-\lambda^* \rho_{g^*}(x, u)] \exp[\lambda^* \rho_{g^*}(y, x)] d\mu_{g^*}(u). \\ &= \exp[\lambda^* \rho_{g^*}(y, x)] \times \int_{\mathcal{M}} \exp[-\lambda^* \rho_{g^*}(x, u)] d\mu_{g^*}(u). \\ &= \exp[\lambda^* \rho_{g^*}(y, x)] \times C(x), \end{aligned}$$

Thus, $\frac{C(y)}{C(x)} \leq \exp[\lambda^* \rho_{g^*}(y, x)]$, which in turn implies $\log \left[\frac{C(y)}{C(x)} \right] \leq \lambda^* \rho_{g^*}(y, x)$. So we have (24) $\leq \lambda^* \rho_{g^*}(x, y) + \lambda^* \rho_{g^*}(y, x)$. Because $\rho_{g^*}(x, y) = \rho_{g^*}(y, x)$, so we have:

$$L_{D, D'}(z) \leq 2\lambda^* \rho_{g^*}(x, y).$$

When $\Delta^* = \sup_{x=\eta(D), y=\eta(D')} \rho_{g^*}(x, y)$, then $L_{D, D'}(z) \leq 2\lambda^* \Delta^*$, Δ^* is the worst-case distance between outputs of adjacent datasets under g^* . Thus, as long as we choose

$$2\lambda^* \Delta^* \leq \varepsilon \implies \lambda^* \leq \frac{\varepsilon}{2\Delta^*} \quad (32)$$

we can prove that $L_{D, D'}(z) \leq \varepsilon$, $\forall z \in \mathcal{M}$, thus we get:

$$\begin{aligned} \Pr[\mathcal{A}(D) \subseteq \mathcal{M}] &= \int_{\mathcal{M}} \mathbb{P}^*(z | \eta(D)) d\mu_{g^*}(z) \\ &\leq e^\varepsilon \int_{\mathcal{M}} \mathbb{P}^*(z | \eta(D')) d\mu_{g^*}(z) \\ &= e^\varepsilon \Pr[\mathcal{A}(D') \subseteq \mathcal{M}]. \end{aligned} \quad (33)$$

which proves that the mechanism is locally ε -DP under the conformal metric g^* . \square

F. Proof for Corollary IV.5

Proof. We prove both directions of the claim, using the definition of ε -DP [9] and triangle inequality for ρ_{g^*} and integration with respect to the Riemannian volume μ_{g^*} .

Sufficiency. Suppose that for all adjacent $x \sim x'$ and all $z \in \mathcal{M}$, the privacy loss is bounded by ε , this means for every outcome z , we have:

$$e^{-\varepsilon} \leq \frac{\mathbb{P}^*(z | x, \lambda^*)}{\mathbb{P}^*(z | x', \lambda^*)} \leq e^\varepsilon.$$

Because the probability densities are nonnegative, we can multiply both sides of the inequality by $\mathbb{P}^*(z | x', \lambda^*)$ to get the pointwise bound for all z :

$$e^{-\varepsilon} \mathbb{P}^*(z | x', \lambda^*) \leq \mathbb{P}^*(z | x, \lambda^*) \leq e^\varepsilon \mathbb{P}^*(z | x', \lambda^*)$$

Now integrate this inequality over an arbitrary measurable set $S \subseteq \mathcal{M}$ with respect to the Riemannian volume measure $d\mu_{g^*}(z)$ (under which $\mathbb{P}^*(\cdot | x)$ is defined as a density). Using the monotonicity of integrals, we obtain:

$$\begin{aligned} e^{-\varepsilon} \int_S \mathbb{P}^*(z | x', \lambda^*) d\mu_{g^*}(z) &\leq \int_S \mathbb{P}^*(z | x, \lambda^*) d\mu_{g^*}(z) \\ &\leq e^\varepsilon \int_S \mathbb{P}^*(z | x', \lambda^*) d\mu_{g^*}(z). \end{aligned}$$

But $\int_S \mathbb{P}^*(z | x, \lambda^*) d\mu_{g^*}(z)$ is exactly $\Pr[\mathcal{A}(D) \in S]$, the probability that mechanism \mathcal{A} outputs an outcome in S given input dataset D (with summary x). Similarly, the rightmost integral is $\Pr[\mathcal{A}(D') \in S]$ for adjacent D' . Therefore, the above inequality is represented as:

$$e^{-\varepsilon} \Pr[\mathcal{A}(D') \in S] \leq \Pr[\mathcal{A}(D) \in S] \leq e^\varepsilon \Pr[\mathcal{A}(D') \in S],$$

for all measurable S and all adjacent D, D' . Thus, \mathcal{A} satisfies ε -DP.

Necessity. Now assume conversely that the mechanism \mathcal{A} is ε -DP. By definition, for any two adjacent datasets D, D' with summaries $x = \eta(D)$ and $x' = \eta(D')$, for any output $z \in \mathcal{M}$, the privacy loss function satisfies: $|\ell_{\mathcal{A}, x, x'}(z)| \leq \varepsilon$. Recall that under (\mathcal{M}, g^*) , the density function is:

$$\mathbb{P}^*(z | x, \lambda^*) = \frac{1}{C(x, \lambda^*)} \exp(-\lambda^* \rho_{g^*}(x, z)),$$

where $C(x, \lambda^*)$ is the normalizing constant

$$C(x, \lambda^*) = \int_{\mathcal{M}} \exp(-\lambda^* \rho_{g^*}(x, u)) d\mu_{g^*}(u).$$

Thus, the privacy loss $\ell_{\mathcal{A}, x, x'}(z)$ can be expanded as:

$$\ell_{\mathcal{A}, x, x'}(z) = \ln \left(\frac{C(x', \lambda^*)}{C(x, \lambda^*)} \right) - \lambda^* (\rho_{g^*}(x, z) - \rho_{g^*}(x', z)).$$

Since ρ_{g^*} is a geodesic distance, it satisfies the triangle inequality:

$$\rho_{g^*}(x, z) \leq \rho_{g^*}(x, x') + \rho_{g^*}(x', z).$$

Rearranging gives:

$$\rho_{g^*}(x, z) - \rho_{g^*}(x', z) \leq \rho_{g^*}(x, x').$$

Similarly, swapping x and x' yields:

$$\rho_{g^*}(x', z) - \rho_{g^*}(x, z) \leq \rho_{g^*}(x, x').$$

Thus, in absolute value, $|\rho_{g^*}(x, z) - \rho_{g^*}(x', z)| \leq \rho_{g^*}(x, x')$. This is a key control for the second term in $\ell_{\mathcal{A}, x, x'}(z)$. Therefore, we obtain the bound:

$$|-\lambda^* (\rho_{g^*}(x, z) - \rho_{g^*}(x', z))| \leq \lambda^* \rho_{g^*}(x, x').$$

We now analyze the normalizing constants $C(x, \lambda^*)$ and $C(x', \lambda^*)$, apply the triangle inequality again:

$$\rho_{g^*}(x', u) \leq \rho_{g^*}(x, u) + \rho_{g^*}(x, x').$$

Thus,

$$\begin{aligned} \exp(-\lambda^* \rho_{g^*}(x', u)) &\geq \exp(-\lambda^* (\rho_{g^*}(x, u) + \rho_{g^*}(x, x'))) \\ &= \exp(-\lambda^* \rho_{g^*}(x, u)) \exp(-\lambda^* \rho_{g^*}(x, x')) \end{aligned}$$

Integrating both sides over $u \in \mathcal{M}$ gives:

$$C(x', \lambda^*) \geq \exp(-\lambda^* \rho_{g^*}(x, x')) C(x, \lambda^*).$$

Thus:

$$\begin{aligned} \frac{C(x', \lambda^*)}{C(x, \lambda^*)} &\geq \exp(-\lambda^* \rho_{g^*}(x, x')) \\ \ln \left(\frac{C(x', \lambda^*)}{C(x, \lambda^*)} \right) &\geq -\lambda^* \rho_{g^*}(x, x') \end{aligned}$$

Finally, we can deduce that the normalizing constant term satisfies:

$$\left| \ln \left(\frac{C(x', \lambda^*)}{C(x, \lambda^*)} \right) \right| \leq \lambda^* \rho_{g^*}(x, x').$$

Combining the above steps, we conclude that the total privacy loss $\ell_{\mathcal{A}, x, x'}(z)$ satisfies:

$$\begin{aligned} |\ell_{\mathcal{A}, x, x'}(z)| &\leq \left| \ln \left(\frac{C(x', \lambda^*)}{C(x, \lambda^*)} \right) \right| + \lambda^* |\rho_{g^*}(x, z) - \rho_{g^*}(x', z)| \\ &\leq 2\lambda^* \rho_{g^*}(x, x') \end{aligned}$$

Thus, if the mechanism \mathcal{A} satisfies ε -DP, it is necessary that

$$2\lambda^* \rho_{g^*}(x, x') \leq \varepsilon,$$

for all adjacent $x \sim x'$. Since x and x' are arbitrary adjacent summaries, this shows that the privacy loss at every output z must satisfy $|\ell_{\mathcal{A}, x, x'}(z)| \leq \varepsilon$. This concludes the necessity proof.

Combining the two directions, we have shown that \mathcal{A} is ε -DP if and only if $|\ell_{\mathcal{A}, x, x'}(z)| \leq \varepsilon$ for all adjacent $x \sim x'$ and all outcomes $z \in \mathcal{M}$. \square

G. Proof for Theorem VI

Proof. Let Theorem IV.4 hold, D and D' are neighboring datasets. Since all data points fall within a geodesic ball $B_r(\mathcal{M})$ of radius r , dimension d , replacing a single sample shifts Fréchet's mean by a distance of up to about $\frac{2r}{n}$ (under the original metric g), which means the global sensitivity $\Delta \leq \frac{2r}{n}$. Since the conformal factor ϕ is bounded (refer to Theorem III.4), we have for any two points $x, y \in B_r(\mathcal{M})$, $\rho_{g^*}(x, y)$ is proportional to $\rho_g(x, y)$, satisfying the inequality (12). Thus $\eta(D)$ also has an upper bound on the sensitivity Δ^* under g^* :

$$\rho_{g^*}(\eta(D), \eta(D')) \leq \sqrt{\phi_{\max}} \rho_g(\eta(D), \eta(D')) \leq \frac{2r\sqrt{\phi_{\max}}}{n}.$$

Thus the rate parameter $\lambda^* = \frac{\varepsilon}{2\Delta^*} \geq \frac{\varepsilon n}{4r\sqrt{\phi_{\max}}}$. Since the output density distribution $\mathbb{P}^*(z|\eta(D))$ is only related to $\rho_{g^*}(z, \eta(D))$, the mechanism is spherically symmetric about $\eta(D)$ under the g^* metric. To calculate the error of the expectation, we switch ρ_g^2 to the conformal metric g^* . Since $\rho_{g^*}(z, \eta) \geq \sqrt{\phi_{\min}} \rho_g(z, \eta)$, for an arbitrarily output z , we have:

$$\rho_g^2(z, \eta(D)) \leq \frac{1}{\phi_{\min}} \rho_{g^*}^2(z, \eta(D)).$$

Thus, the upper bound of the expectation error is:

$$\begin{aligned} \mathbb{E}[\rho_g^2(\eta(D), \mathcal{A}(D))] &= \int_{\mathcal{M}} \rho_g^2(z, \eta(D)) \cdot \mathbb{P}^*(z | \eta(D)) dz \\ &< \frac{1}{\phi_{\min}} \int_{\mathcal{M}} \rho_{g^*}^2(z, \eta(D)) \cdot \mathbb{P}^*(z | \eta(D)) dz \\ &= \frac{1}{\phi_{\min}} \mathbb{E}[\rho_{g^*}^2(\eta(D), \mathcal{A}(D))] \end{aligned}$$

Now, we need to calculate the explicit form of $\mathbb{E}[\rho_{g^*}^2(\eta(D), \mathcal{A}(D))]$. Let random variable $T = \rho_{g^*}(\eta(D), \mathcal{A}(D))$ represents the distance of the mechanism's output relative to $\eta(D)$ under g^* . Due to the spherical symmetry, the probability density function of T can be expressed in polar coordinate form as:

$$PDF_T(t) = Ct^{d-1} e^{-\lambda^* t}, \quad t \geq 0$$

where the normalization condition determines the constant C . By utilizing Gamma function [75],

$$\int_0^\infty t^{d-1} e^{-\lambda^* t} dt = \frac{\Gamma(d)}{(\lambda^*)^d} = \frac{(d-1)!}{(\lambda^*)^d}.$$

Thus, $C = \frac{(\lambda^*)^d}{(d-1)!}$, and $PDF_T(t) = \frac{(\lambda^*)^d}{(d-1)!} t^{d-1} e^{-\lambda^* t}$, $t \geq 0$. Based on the PDF_T , we calculate the second-order variance:

$$\mathbb{E}[T^2] = \int_0^\infty t^2 p_T(t) dt = \frac{(\lambda^*)^d}{(d-1)!} \int_0^\infty t^{d+1} e^{-\lambda^* t} dt.$$

Combining with Gamma function $\int_0^\infty t^{d+1} e^{-\lambda^* t} dt = \frac{(d+1)!}{(\lambda^*)^{d+2}}$, we have:

$$\mathbb{E}[T^2] = \frac{(\lambda^*)^d}{(d-1)!} \cdot \frac{(d+1)!}{(\lambda^*)^{d+2}} = \frac{(d+1)!}{(d-1)!} \cdot \frac{1}{(\lambda^*)^2} = \frac{d(d+1)}{(\lambda^*)^2}.$$

Combining with $\lambda^* = \frac{\varepsilon}{2\Delta^*} \geq \frac{\varepsilon n}{4r\sqrt{\phi_{\max}}}$ (worst-case), we have the upper bound of expected square error under metric g as follows:

$$\begin{aligned} \mathbb{E}[\rho_g^2(\eta(D), \mathcal{A}(D))] &< \frac{1}{\phi_{\min}} \frac{d(d+1)}{(\lambda^*)^2} \\ &= \frac{d(d+1)}{\phi_{\min}} \frac{16r^2 \phi_{\max}}{\varepsilon^2 n^2} \\ &= \frac{16d(d+1)r^2 \phi_{\max}}{\varepsilon^2 n^2 \phi_{\min}} \end{aligned}$$

This completes the proof of Theorem V.1. \square

H. Implementation Details for Section. VII-A

We provide further implementation details and hyperparameters of our proposed Conformal Differential Privacy and Riemannian-Laplace DP by Reimherr et al. [18] and tangent Gaussian DP by Utpala et al. [19] as follows:

Hyperparameters. We test for algorithm utility under privacy budgets $\varepsilon \in \{0.1, 0.2, \dots, 0.9\}$ and $\delta = 10^{-9}$. CIFAR-10 [67] contains a total of 60,000 images in 10 classes, with 6,000 images for each class; Fashion-MNIST [68] has 70,000 images in total, with 7,000 images for each class. Our selected $\delta = 10^{-9}$ satisfies $\delta \ll \frac{1}{|D|}$ for both datasets. The *global* sensitivity Δ is calculated as $\Delta = \frac{2r}{N}$, in which r is the radius of the geodesic ball containing all images, N is the number of data samples in the dataset. For more calculations in determining r , we refer to the work of Utpala et al. [19].

Implementation. We implement both Riemannian-Laplace DP [18] and tangent Gaussian DP [19] with the same parameters described in Utpala et al. [19], specifically, for the global sensitivity of $\Delta_{glb} = \frac{2r}{N}$, in which r is the radius of the geodesic ball that contains all data samples. We adopt the *classical* Gaussian noise described by Utpala et al. [19], in which we have $\sigma = \frac{\Delta_{glb}}{\varepsilon} \sqrt{2 \ln(\frac{1.25}{\delta})}$. The algorithms and calculations over the manifold are implemented with the `geomstats` Python library by Miolane et al. [76]. The calculation of the Fréchet mean on the manifold through gradient descent can be easily implemented with `geomstats`. We thank the authors for the invaluable help in making the code publicly available.

The MCMC process utilized by both our proposed *Conformal-DP* mechanism is shown in Algorithm 1.

Algorithm 1 Differentially Private Conformal Laplace Sampling on Manifolds

- 1: **Input:**
 - 2: Conformal Fréchet mean $\eta(D)$,
 - 3: Privacy budget $\varepsilon > 0$,
 - 4: Local sensitivity bound $\Delta^*(D)$ (estimated or numerical),
 - 5: Sampling radius r ($r = \infty$ for global sampling),
 - 6: Proposal distribution $Q(\cdot)$ (MCMC),
 - 7: Maximum iterations T_{\max}
 - 8: **Output:**
 - 9: Differentially private sample Z from $\mathbb{P}_r^*(z | \eta(D), \lambda^*)$
 - 10: **Step 1: Calculate Noise Intensity**
 - $\lambda^* \leftarrow \min \left\{ \frac{\varepsilon}{2\Delta^*(D)}, \lambda_{\max} \right\}$ (Default: $\lambda^* = \frac{\varepsilon}{2\Delta^*(D)}$)
 - 11: **Step 2: Select Sampling Method**
 - 12: **if MH [77] then**
 - 13: Initialize $z^{(0)} \in B_r^*(\eta(D))$ ▷ Starting point
 - 14: **for** $t = 0$ **to** $T_{\max} - 1$ **do**
 - 15: Propose $z' \sim Q(\cdot | z^{(t)})$ ▷ Generate candidate
 - 16: Compute acceptance ratio:
 - $\alpha \leftarrow \frac{\exp[-\lambda^* \rho_{g'}(\eta(D), z')]}{\exp[-\lambda^* \rho_{g'}(\eta(D), z^{(t)})]} \frac{Q(z^{(t)} | z')}{Q(z' | z^{(t)})}$
 - 17: Accept $z^{(t+1)} \leftarrow z'$ with probability $\min(1, \alpha)$
 - 18: **end for**
 - 19: Set $Z \leftarrow z^{(T_{\max})}$ ▷ Final sample
 - 20: **else** ▷ Semi-analytic method for specific manifolds
 - 21: Exploit manifold structure (e.g., sphere, hyperbolic space)
 - 22: Directly sample from $\mathbb{P}_r^*(z | \eta(D), \lambda^*)$
 - 23: **end if**
 - 24: **Return:** Z ▷ ε -DP perturbed output
-