Electronic Research Archive

https://www.aimspress.com/journal/era

ERA, 34(9): 6043–6071.
DOI: 10.3934/era.2026267
Received: 19 March 2026
Revised: 15 June 2026
Accepted: 26 June 2026
Published: 14 July 2026

Research article

Operator normed adaptive Laplace mechanism: A manifold-based approach for differential privacy

Bekir Danış*

Department of Computer Engineering, Faculty of Engineering, Aydın Adnan Menderes University, Aydın 09010, Türkiye

* **Correspondence:** Email: bekir.danis@adu.edu.tr.

Abstract: Modern computational intelligence and data science applications frequently operate on complex, non-Euclidean domains. Applying differential privacy (DP) to these structures requires mechanisms that respect their underlying geometry. Instance-adaptive mechanisms in DP improve utility over global worst-case bounds but often require heavy-tailed noise, leading to numerical instability and unbounded variance. We introduced a geometric framework that lifted sensitivity analysis to the tangent bundle of a Riemannian manifold. The proposed directional local sensitivity operator captured maximal directional derivatives along minimizing geodesics, enabling our operator normed adaptive Laplace mechanism (ON–ALM). By enforcing a β-log-Lipschitz stability condition on the regularized sensitivity, noise was precisely calibrated via directional operator norms. Under (ε, δ)-DP, extracting a typical set allowed the use of light-tailed Laplace noise, achieving strictly bounded variance, improved numerical stability, and high utility for geometric and directional queries.

Keywords: differential privacy; geometric sensitivity; operator norm; manifold-based modelling; Laplace mechanisms; computational intelligence; next-generation data science

1. Introduction

Differential privacy (DP) [1] provides a rigorous probabilistic framework for safeguarding individual-level information in statistical analyses. In the era of next-generation data science, computational intelligence models (such as deep representation learning and natural language embeddings) frequently operate on complex, non-Euclidean domains modeled as Riemannian manifolds. Securing these advanced learning algorithms requires privacy mechanisms that respect the underlying geometry of the data. By ensuring that the output distribution of a randomized algorithm remains nearly unchanged when a single data record is modified, DP offers quantifiable protection against reconstruction and inference attacks. The magnitude of noise required for privacy is determined by the sensitivity of the query, the maximal change in output between neighboring datasets.

Classically, noise is calibrated to the global sensitivity of a query [1]. While this worst-case approach guarantees strong privacy, it often results in excessive noise, especially for queries exhibiting moderate variation on typical datasets. To address this conservatism, data-dependent notions such as local sensitivity have been proposed. However, directly using local sensitivity to calibrate noise can violate differential privacy, as sensitivity itself may vary sharply across adjacent datasets. The smooth sensitivity framework [2] mitigates this issue by regularizing local sensitivity over the data space. In the pure ε-DP setting, such constructions often require heavy-tailed noise to uniformly control privacy loss, which may lead to infinite variance and numerical instability.

Recent work aims to integrate instance-adaptive calibration with light-tailed noise. Mechanisms such as the piecewise Laplace [3] achieve pure DP through refined structural arguments, while Gaussian DP [4] provides alternative (ε, δ)-style guarantees via hypothesis-testing characterizations. Geometric perspectives have also begun shaping mechanism design: for example, Reimherr and Awan [5] introduced the K-norm gradient mechanism, aligning noise with query geometry through norm-based constructions.

Recent studies have increasingly investigated differential privacy in non-Euclidean and manifold-valued settings. Han et al. proposed a differentially private Riemannian optimization framework based on intrinsic Gaussian perturbations on tangent spaces [6]. Long et al. combined manifold optimization and differential privacy in mobile crowd sensing applications [7]. Huang et al. introduced a federated learning framework on Riemannian manifolds with formal DP guarantees [8]. Moreover, privacy-preserving mechanisms on manifolds have been investigated using Fréchet means and geodesic distance structures [9]. Very recently, conformal transformations have been employed to improve privacy-utility trade-offs on heterogeneous manifold data [10]. These developments indicate the growing importance of geometry-aware privacy mechanisms and motivate the proposed operator normed adaptive Laplace mechanism (ON–ALM) framework.

Motivated by these developments, we adopt a differential-geometric viewpoint and formulate sensitivity intrinsically on a Riemannian manifold. We represent the dataset space via a smooth embedding into a complete Riemannian manifold and interpret adjacency through bounded geodesic distance. Within this setting, first-order query variation is captured by the differential of the mapping, and we define the *Directional Local Sensitivity Operator* Δ^* as the operator norm of df on the tangent space. This formulation quantifies the maximal directional variation of the query and provides a geometrically structured refinement of classical sensitivity.

Building on this operator-theoretic representation, we introduce the ON–ALM, which calibrates multivariate Laplace noise to a regularized version of Δ^* satisfying a β-log-Lipschitz stability condition. Under this assumption, explicit noise scales are derived, and rigorous (ε, δ)-DP guarantees are established via a typical-set-based privacy-loss analysis. Unlike some pure ε-DP adaptive schemes that require heavy-tailed noise, our construction demonstrates that light-tailed Laplace perturbations with strictly bounded variance suffice for stable, instance-adaptive calibration.

Beyond formal privacy guarantees, the framework reveals the geometric structure underlying sensitivity variation. The stability parameter β can be interpreted in terms of the second-order regularity of the query along geodesics, connecting privacy calibration to intrinsic smoothness properties. This perspective provides a principled design rule: measure sensitivity via directional operator norms on the tangent bundle, and align noise geometry with this structure through ℓ_1-calibrated Laplace perturbations.

In summary, our contributions are twofold. First, we provide a geometrically grounded, operator-theoretic formulation of local sensitivity on manifolds. Second, we show that, under mild regularity conditions, this formulation enables stable, finite-variance, instance-adaptive mechanisms satisfying (ε, δ)-DP. This approach establishes a transparent and mathematically rigorous pathway for geometry-aware privacy mechanisms suitable for directional, structured, and manifold-valued queries.

Positioning relative to related work. Table 2 in Section 5 summarizes how ON–ALM relates to the primary competing mechanisms. Briefly: the classical Laplace mechanism [1] uses global sensitivity and is oblivious to local geometry; the smooth sensitivity framework [2] achieves instance-adaptivity but requires heavy-tailed noise and a global optimization over the data space; the K-norm gradient mechanism [5] aligns noise geometry with query structure via convex bodies but operates in Euclidean output spaces without intrinsic manifold sensitivity; Gaussian DP [4] provides (ε, δ)-guarantees with light-tailed noise but does not exploit local sensitivity variation. ON–ALM occupies the intersection of instance-adaptivity, light-tailed Laplace noise, and intrinsic Riemannian geometry, at the cost of a feasibility constraint on β that limits applicability to low- or moderate-dimensional outputs. The precise regime of advantage is characterized in Section 5.

2. Preliminaries: Geometric embedding of datasets and DP

In this section, we briefly review the essential concepts of Riemannian geometry used throughout the paper. For a comprehensive treatment and detailed proofs, we refer the reader to standard texts such as [11, 12].

2.1. Discrete dataset space and Hamming geometry

Let $\mathcal{D} = \mathcal{X}^N$ denote the space of datasets of size N, where each record lies in a data universe $\mathcal{X}$. We equip $\mathcal{D}$ with the Hamming metric

$$d_H(D, D') = \#\{i \in \{1, \dots, N\} : x_i \neq x_i'\}.$$

Two datasets $D, D' \in \mathcal{D}$ are said to be *adjacent*, denoted $D \sim D'$, if $d_H(D, D') = 1$.

2.2. Geometric embedding into a Riemannian manifold

To enable differential analysis, we embed the discrete dataset space into a smooth ambient space. Let

$$\Phi : \mathcal{D} \to (M, g)$$

be an injective mapping into a smooth, complete Riemannian manifold (M, g).

Assumption 1 (Controlled Geometric Embedding)**.** *The embedding Φ satisfies the following properties:*

1) ***Injectivity:*** *Φ is injective.*

2) ***Adjacency Lipschitz Control:*** *There exists a constant $L > 0$ such that for all adjacent datasets $D \sim D'$,*

$$d_M(\Phi(D), \Phi(D')) \leq L.$$

3) ***Local Geodesic Regularity:*** *The injectivity radius of M along $\Phi(\mathcal{D})$ satisfies*

$$\inf_{x\in\Phi(\mathcal{D})} \mathrm{inj}(x) > L.$$

By rescaling the metric $g \mapsto g/L^2$, we may assume without loss of generality that $L = 1$, hence:

$$d_M(\Phi(D), \Phi(D')) \leq 1 \quad \text{whenever } D \sim D'.$$

This normalization preserves geodesic structure and ensures that adjacency corresponds to a bounded geodesic displacement.

For any two adjacent datasets $x, x' \in M$, we have $d_M(x, x') \leq 1$. By our assumption that the injectivity radius satisfies $\mathrm{inj}(M) > 1$, the point x' lies strictly inside the injectivity domain of x, and thus entirely outside its cut locus. Consequently, x and x' are connected by a *unique* minimizing unit-speed geodesic γ that depends smoothly on its endpoints. This guarantees that the differential operations, integration along γ, and parallel transport are rigorously well-defined and free of singularities.

2.3. Canonical product-manifold construction

A fundamental example arises when $\mathcal{X}$ is itself a compact Riemannian manifold. Consider a dataset

$$D = (x_1, \ldots, x_N), \quad x_i \in \mathcal{X},$$

and define

$$M = \mathcal{X}^N$$

equipped with the product Riemannian metric

$$g_M = \sum_{i=1}^{N} g_{\mathcal{X}}^{(i)}.$$

Lemma 1 (Product-Manifold Embedding)**.** *Let $\Phi(D) = (x_1, \ldots, x_N) \in M$. If $D \sim D'$ differ in exactly one coordinate j, then*

$$d_M(\Phi(D), \Phi(D')) = d_{\mathcal{X}}(x_j, x'_j) \leq \mathrm{diam}(\mathcal{X}),$$

where $\mathrm{diam}(\mathcal{X}) := \sup_{x,y\in\mathcal{X}} d_{\mathcal{X}}(x, y) < \infty$.

Proof. Under the product metric,

$$d_M^2(x, y) = \sum_{i=1}^{N} d_{\mathcal{X}}(x_i, y_i)^2.$$

If $x_i = y_i$ for all $i \neq j$, only one term remains, yielding the result.

2.4. Example: Product manifold on the sphere

A canonical example arises when individual records lie on the unit sphere $\mathcal{X} = \mathbb{S}^{d-1}$. A dataset of size N is then embedded in

$$M = (\mathbb{S}^{d-1})^N,$$

with the product Riemannian metric. For adjacent datasets differing in the i-th coordinate $u_i \neq u_i'$, the geodesic distance satisfies

$$d_M(\Phi(D), \Phi(D')) = d_{\mathbb{S}^{d-1}}(u_i, u_i').$$

If the angular variation is bounded, adjacency corresponds to a bounded geodesic displacement in M, and discrete perturbations are naturally represented as unit-speed minimizing geodesics along the i-th factor.

2.5. Practical construction of the embedding Φ

A natural question arising in the application of the geometric framework is how to construct a concrete embedding $\Phi : \mathcal{D} \to (M, g)$ for a real dataset. We present two representative examples that cover a broad range of practical scenarios.

Example 1 (Mean Embedding for Tabular Datasets)**.** *Let $\mathcal{X} = \mathbb{R}^d$ and consider datasets $D = (x_1, \dots, x_N) \in (\mathbb{R}^d)^N$. A canonical embedding is the empirical mean map*

$$\Phi(D) = \frac{1}{N}\sum_{i=1}^{N} x_i \ \in \ \mathbb{R}^d,$$

with the ambient space $(M, g) = (\mathbb{R}^d, g_{\text{euc}})$ equipped with the standard Euclidean metric. For adjacent datasets $D \sim D'$ differing in one record $x_j \mapsto x_j'$,

$$d_M(\Phi(D), \Phi(D')) = \frac{1}{N}\|x_j - x_j'\|_2 \leq \frac{B}{N},$$

where $B = \sup_{x,x'\in\mathcal{X}} \|x - x'\|_2$ is the diameter of the data universe. Rescaling by N/B achieves $L = 1$ in Assumption 1. Since $\mathbb{R}^d$ is a flat, simply connected, complete Riemannian manifold with infinite injectivity radius, all conditions of Assumption 1 are satisfied. For a scalar-valued query $f(\Phi(D)) = \langle w, \Phi(D)\rangle$, one obtains $\Delta^* f \equiv \|w\|_2$, recovering the classical ℓ_2-sensitivity of the mean. This embedding is immediately applicable to standard tabular datasets where each record is a d-dimensional feature vector, such as census microdata or medical records.

Example 2 (Normalized Embedding into the Positive Orthant via the symmetric positive definite (SPD) Manifold)**.** *Consider a setting where each record $x_i \in \mathbb{R}^d$ is a unit-norm feature vector ($\|x_i\|_2 = 1$), and the dataset summary of interest is the sample covariance matrix*

$$\Phi(D) = \frac{1}{N}\sum_{i=1}^{N} x_i x_i^\top \ \in \ \mathcal{S}_{++}^d,$$

where $\mathcal{S}_{++}^d$ denotes the manifold of $d \times d$ SPD matrices. We equip $\mathcal{S}_{++}^d$ with the affine-invariant Riemannian metric

$$g_P(U, V) = \operatorname{tr}(P^{-1}UP^{-1}V), \quad U, V \in T_P\mathcal{S}_{++}^d \cong \mathcal{S}^d,$$

where $\mathcal{S}^d$ is the space of $d \times d$ symmetric matrices. This manifold is geodesically complete with nonpositive sectional curvature, and the injectivity radius satisfies $\mathrm{inj}(P) = \infty$ *for all P. For adjacent datasets $D \sim D'$ differing in the j-th record,*

$$\Phi(D') - \Phi(D) = \frac{1}{N}(x'_j(x'_j)^\top - x_j x_j^\top),$$

and the geodesic distance $d_{\mathcal{S}^d_{++}}(\Phi(D), \Phi(D'))$ is bounded uniformly in terms of $1/N$ *and* $\|x_j\|_2 = \|x'_j\|_2 = 1$*. This embedding is directly applicable to covariance-based queries arising in brain-computer interface (BCI) data analysis [11], diffusion tensor imaging, and portfolio-risk estimation, where the underlying data naturally resides on the SPD manifold.*

2.6. Differential structure and operator norm

Let $f : M \to \mathbb{R}^k$ be a C^1 function. For each $x \in M$, the differential

$$df_x : T_xM \to \mathbb{R}^k$$

is a bounded linear operator. The tangent space T_xM is equipped with the Riemannian norm $\|\cdot\|_g$, and $\mathbb{R}^k$ with the ℓ_1 norm.

We define the operator norm

$$\|df_x\|_{\mathrm{op}} := \sup_{\|v\|_g=1} \|df_x(v)\|_1.$$

This choice naturally reflects the ℓ_1 geometry of multivariate Laplace perturbations.

Remark 1. *The operator norm defined above is induced by the Riemannian norm on the tangent space T_xM and the ℓ_1 norm on $\mathbb{R}^k$, hence it is a mixed operator norm.*

For scalar-valued queries ($k = 1$), the codomain is equipped with the absolute value, and the operator norm coincides with the dual norm on T_x^*M. Using the metric-induced isometric isomorphism between T_xM and T_x^*M, we obtain

$$\|df_x\|_{\mathrm{op}} = \|df_x\|_{g,*} = \|\nabla f(x)\|_g.$$

For multivariate queries ($k > 1$), writing $f = (f_1, \dots, f_k)$, we have

$$\|df_x\|_{\mathrm{op}} = \sup_{\|v\|_g=1} \sum_{i=1}^{k} |df_i(x)(v)|.$$

By the triangle inequality and duality,

$$\|df_x\|_{\mathrm{op}} \le \sum_{i=1}^{k} \|\nabla f_i(x)\|_g.$$

In general, this bound is not tight; the operator norm captures the maximal joint directional variation across all components. This formulation provides a natural geometric generalization of gradient-based sensitivity to vector-valued queries on manifolds.

2.7. Mechanisms and differential privacy

We consider mechanisms of the form

$$\mathcal{M}(D) = f(\Phi(D)) + Z,$$

where $Z \in \mathbb{R}^k$ is a random vector.

Definition 1 (DP)**.** *A mechanism $\mathcal{M}$ satisfies (ε, δ)-DP if for all $D \sim D'$ and all measurable sets $S \subset \mathbb{R}^k$,*

$$\mathbb{P}[\mathcal{M}(D) \in S] \leq e^{\varepsilon}\mathbb{P}[\mathcal{M}(D') \in S] + \delta.$$

2.8. Geometric sensitivity via geodesic interpolation

The embedding enables a continuous interpolation between adjacent datasets. Let $D \sim D'$ and let $\gamma : [0, \ell] \to M$ be a minimizing unit-speed geodesic connecting $\Phi(D)$ to $\Phi(D')$, where $\ell = d_M(\Phi(D), \Phi(D')) \leq 1$.

Proposition 1 (Geodesic Sensitivity Bound)**.** *Let $f : M \to \mathbb{R}^k$ be C^1. Then for all adjacent $D \sim D'$,*

$$\|f(\Phi(D)) - f(\Phi(D'))\|_1 \leq \int_0^{\ell} \|df_{\gamma(t)}\|_{\mathrm{op}}\, dt \leq \ell \sup_{x \in M} \|df_x\|_{\mathrm{op}}.$$

In particular, under the normalized embedding ($\ell \leq 1$),

$$\Delta f \leq \sup_{x \in M} \|df_x\|_{\mathrm{op}}.$$

Proof. Applying the fundamental theorem of calculus along γ,

$$f(\Phi(D')) - f(\Phi(D)) = \int_0^{\ell} df_{\gamma(t)}(\dot{\gamma}(t))\, dt.$$

Taking ℓ_1 norms and using $\|\dot{\gamma}(t)\|_g = 1$ yields

$$\|df_{\gamma(t)}(\dot{\gamma}(t))\|_1 \leq \|df_{\gamma(t)}\|_{\mathrm{op}},$$

which proves the result.

Remark 2 (Global versus Local Sensitivity Regimes)**.** *The bound*

$$\Delta f \leq \sup_{x \in M} \|df_x\|_{\mathrm{op}}$$

provides a worst-case global sensitivity estimate. However, if M is non-compact or if $\|df_x\|_{\mathrm{op}}$ is not uniformly bounded, this supremum may be infinite.

In contrast, for each fixed $x \in M$, the differential df_x is a bounded linear operator, so the pointwise quantity $\|df_x\|_{\mathrm{op}}$ is always finite. This distinction motivates the transition from global to local sensitivity.

The framework developed in this work replaces the global supremum with a regularized, pointwise sensitivity functional satisfying a log-Lipschitz stability condition. This allows finite noise calibration even in settings where classical global sensitivity is not well-defined.

3. Directional local sensitivity operator

This section develops a fully intrinsic, coordinate-invariant formulation of local sensitivity on Riemannian manifolds. All geometric assumptions are stated explicitly, and regularity properties of the proposed sensitivity functional are established with complete analytical rigor.

Throughout this section we operate under the geometric embedding conditions introduced in Assumption 1.

Remark 3. *Assumption 1 ensures that adjacent datasets are connected by a minimizing unit-speed geodesic. The injectivity radius condition prevents cut-locus phenomena and guarantees smooth dependence of geodesics on endpoints.*

3.1. Definition and regularity of the sensitivity operator

Let $f \in C^1(M, \mathbb{R}^k)$. For each $x \in M$, the differential

$$df_x : T_xM \to \mathbb{R}^k$$

is a bounded linear operator.

We equip T_xM with the norm induced by the Riemannian metric g, and $\mathbb{R}^k$ with the ℓ_1 norm. The induced operator norm is

$$\|df_x\|_{\mathrm{op}} = \sup_{\|v\|_g=1} \|df_x(v)\|_1.$$

Definition 2 (Directional Local Sensitivity Operator)**.** *The directional local sensitivity operator of f is the function*

$$\Delta^* f : M \to [0, \infty), \qquad \Delta^* f(x) = \|df_x\|_{\mathrm{op}}.$$

Proposition 2 (Continuity)**.** *If $f \in C^1(M, \mathbb{R}^k)$, then $\Delta^* f$ is continuous on M.*

Proof. Since $f \in C^1$, the map $x \mapsto df_x$ is continuous as a section of the bundle $\mathrm{Hom}(TM, \mathbb{R}^k)$. The operator norm is continuous on the space of bounded linear maps. Therefore, $x \mapsto \|df_x\|_{\mathrm{op}}$ is continuous as the composition of continuous maps.

3.2. Geodesic sensitivity bound

Proposition 3 (Intrinsic Sensitivity Bound)**.** *Let $f \in C^1(M, \mathbb{R}^k)$. Let $D, D' \in \mathcal{D}$ be adjacent, and denote $x = \Phi(D)$, $x' = \Phi(D')$. Let $\gamma : [0, \ell] \to M$ be a minimizing unit-speed geodesic joining them, where $\ell = d_M(x, x')$. Then,*

$$\|f(x') - f(x)\|_1 \le \int_0^\ell \Delta^* f(\gamma(t))\, dt.$$

In particular,

$$\|f(x') - f(x)\|_1 \le \ell \sup_{t\in[0,\ell]} \Delta^* f(\gamma(t)).$$

Proof. Define $h(t) = f(\gamma(t))$. Since f and γ are smooth, $h \in C^1([0, \ell], \mathbb{R}^k)$. By the fundamental theorem of calculus,

$$f(x') - f(x) = \int_0^\ell h'(t)\, dt.$$

By the chain rule,

$$h'(t) = df_{\gamma(t)}(\dot{\gamma}(t)).$$

Since $\|\dot{\gamma}(t)\|_g = 1$,

$$\|h'(t)\|_1 \leq \Delta^* f(\gamma(t)).$$

Taking norms and integrating yields the result.

Corollary 1 (Adjacency Sensitivity Bound)**.** *Under the normalized geometry (Assumption 1), let $D, D' \in \mathcal{D}$ be adjacent datasets and denote $x = \Phi(D)$, $x' = \Phi(D')$. Let $\gamma : [0, \ell] \to M$ be a minimizing unit-speed geodesic joining them, where $\ell = d_M(x, x') \leq 1$. Then,*

$$\|f(\Phi(D')) - f(\Phi(D))\|_1 \leq \sup_{t\in[0,\ell]} \Delta^* f(\gamma(t)).$$

Proof. From Proposition 2,

$$\|f(x') - f(x)\|_1 \leq \ell \sup_{t\in[0,\ell]} \Delta^* f(\gamma(t)).$$

By our normalized geometric embedding, we have $\ell = d_M(x, x') \leq 1$. Substituting this bound immediately yields the result.

3.3. Log-Lipschitz stability

For $\xi > 0$, define the regularized sensitivity

$$\Delta^*_\xi f(x) = \max\{\Delta^* f(x), \xi\}.$$

Definition 3 (β-Log-Lipschitz Stability)**.** *$\Delta^*_\xi f$ is β-log-Lipschitz if*

$$|\ln \Delta^*_\xi f(x) - \ln \Delta^*_\xi f(y)| \leq \beta\, d_M(x, y) \quad \forall x, y \in M.$$

Remark 4 (Practical Selection of the Regularization Parameter ξ)**.** *The regularization parameter $\xi > 0$ serves as a lower bound on the sensitivity, preventing degeneracy when $\Delta^* f(x) \approx 0$. Its choice has direct consequences for the stability constant $\beta \leq C/\xi$, the noise scale $b(x)$, and the feasibility of the mechanism. We provide three complementary strategies for selecting ξ in practice.*

Strategy 1: Hessian-based analytic bound. *If a global Hessian bound $\|\nabla^2 f(x)\|_{\mathrm{op}} \leq C$ is available analytically or by inspection, choose ξ to enforce a target stability β_0:*

$$\xi = \frac{C}{\beta_0}.$$

For example, with a linear query $f(x) = \langle w, x\rangle$ one has $C = 0$ and any $\xi > 0$ trivially satisfies $\beta = 0$.

Strategy 2: Data-driven estimation. *Evaluate $\Delta^* f$ on a small held-out public dataset $\mathcal{D}_{\mathrm{pub}}$ and set*

$$\xi = \alpha \cdot \min_{x\in\mathcal{D}_{\mathrm{pub}}} \Delta^* f(x),$$

for a safety factor $\alpha \in (0.1, 0.5)$. This ensures that the regularization floor lies below the empirically observed sensitivity minimum, so the max in $\Delta^*_\xi f = \max\{\Delta^* f, \xi\}$ is rarely active, preserving utility.

Strategy 3: Feasibility-backward selection. *Given target privacy parameters (ε, δ) and output dimension k, choose ξ to guarantee that the feasibility condition is satisfied. From $\beta \le C/\xi$ and $\varepsilon > 2k\beta$ (asymptotically), one requires $\xi > 2kC/\varepsilon$. A practical choice is*

$$\xi = \frac{2kC}{\varepsilon} \cdot (1 + \eta),$$

for a small slack $\eta > 0$, e.g., $\eta = 0.1$.

Privacy-utility trade-off. *Increasing ξ reduces β and thereby enlarges the feasible regime, but raises the noise floor $b(x) \ge e^{2\beta}\xi/(\varepsilon - k\beta - \Lambda_{k,\delta}(e^{\beta} - 1))$. Decreasing ξ tightens noise for points where $\Delta^* f(x)$ is large, but may push β above the feasibility threshold. The optimal ξ trades off these effects and in practice can be tuned via a grid search over $\xi \in [10^{-3}, 10^{-1}]$ with validation on a utility metric (e.g., mean squared error on a synthetic release).*

Remark 5. *For a C^2 mapping $f : M \to \mathbb{R}^k$, the differential df is a section of $T^*M \otimes \mathbb{R}^k$. Its covariant derivative with respect to the Levi-Civita connection defines the Riemannian Hessian $\nabla^2 f$, which is a symmetric $(0, 2)$-tensor field taking values in $\mathbb{R}^k$. We define its mixed tensor operator norm at $x \in M$ by taking the supremum over unit tangent vectors:*

$$\|\nabla^2 f(x)\|_{\mathrm{op}} := \sup_{\|u\|_g=1, \|v\|_g=1} \|\nabla^2 f(x)(u, v)\|_1.$$

Proposition 4 (Hessian Control Implies Log-Lipschitz Stability)**.** *Let (M, g) be a complete Riemannian manifold and let $f \in C^2(M, \mathbb{R}^k)$. Assume that*

$$\|\nabla^2 f(x)\|_{\mathrm{op}} \le C \quad \forall x \in M.$$

Then, for every $\xi > 0$, the regularized sensitivity

$$\Delta^*_\xi f(x) := \max\{\|df_x\|_{\mathrm{op}}, \xi\}$$

satisfies the log-Lipschitz estimate

$$\left| \ln \Delta^*_\xi f(x) - \ln \Delta^*_\xi f(y) \right| \le \frac{C}{\xi} d_M(x, y) \quad \forall x, y \in M.$$

Proof. Let $x, y \in M$. If $d_M(x, y) = 0$, the result is trivial, so assume $d_M(x, y) > 0$. By completeness and the injectivity radius condition, there exists a minimizing unit-speed geodesic $\gamma : [0, \ell] \to M$ with $\gamma(0) = x$, $\gamma(\ell) = y$, where $\ell = d_M(x, y)$ and $\|\dot\gamma(t)\|_g = 1$.

Step 1: Reduction to a fixed Banach space via parallel transport.

Let $P_{0,t} : T_x M \to T_{\gamma(t)} M$ denote parallel transport along γ. By metric compatibility, $P_{0,t}$ is a linear isometry. Define

$$A(t) := df_{\gamma(t)} \circ P_{0,t} \in \mathcal{L}(T_x M, \mathbb{R}^k).$$

Step 2: Derivative bound in operator norm.

For $v \in T_xM$ with $\|v\|_g = 1$, define $v_t := P_{0,t}(v)$, so $\|v_t\|_g = 1$. Then,

$$A(t)(v) = df_{\gamma(t)}(v_t).$$

Differentiating along γ and using $\nabla_{\dot{\gamma}(t)} v_t = 0$, we obtain

$$\frac{d}{dt} A(t)(v) = \nabla^2 f(\gamma(t))(\dot{\gamma}(t), v_t).$$

Taking ℓ^1-norms and applying the operator norm bound (using the definition $\|\nabla^2 f\|_{\mathrm{op}} = \sup_{\|u\|_g=\|v\|_g=1} \|\nabla^2 f(u, v)\|_1$ from the preceding Remark):

$$\left\| \frac{d}{dt} A(t)(v) \right\|_1 = \|\nabla^2 f(\gamma(t))(\dot{\gamma}(t), v_t)\|_1 \le \|\nabla^2 f(\gamma(t))\|_{\mathrm{op}} \, \|\dot{\gamma}(t)\|_g \, \|v_t\|_g = \|\nabla^2 f(\gamma(t))\|_{\mathrm{op}} \le C.$$

Taking the supremum over $\|v\|_g = 1$ yields

$$\left\| \frac{d}{dt} A(t) \right\|_{\mathrm{op}} \le C.$$

Step 3: Mean value inequality in Banach space.

Since $A(t)$ takes values in the Banach space $\mathcal{L}(T_xM, \mathbb{R}^k)$, we apply the integral form of the mean value inequality:

$$\|A(\ell) - A(0)\|_{\mathrm{op}} \le \int_0^\ell \left\| \frac{d}{dt} A(t) \right\|_{\mathrm{op}} dt \le C\ell.$$

Step 4: Lipschitz continuity of the operator norm.

We use the fact that the operator norm is 1-Lipschitz on $\mathcal{L}(T_xM, \mathbb{R}^k)$: for any linear operators T, S,

$$\left| \|T\|_{\mathrm{op}} - \|S\|_{\mathrm{op}} \right| \le \|T - S\|_{\mathrm{op}}.$$

(This follows directly from the definition of the operator norm via the triangle inequality.)

Applying this with $T = A(\ell)$ and $S = A(0)$ gives

$$\left| \|A(\ell)\|_{\mathrm{op}} - \|A(0)\|_{\mathrm{op}} \right| \le \|A(\ell) - A(0)\|_{\mathrm{op}} \le C\ell.$$

Step 5: Identification with differentials.

Since parallel transport is an isometry,

$$\|A(\ell)\|_{\mathrm{op}} = \|df_y\|_{\mathrm{op}}, \quad \|A(0)\|_{\mathrm{op}} = \|df_x\|_{\mathrm{op}}.$$

Thus,

$$\left| \|df_y\|_{\mathrm{op}} - \|df_x\|_{\mathrm{op}} \right| \le C d_M(x, y).$$

Step 6: Lipschitz property of the logarithmic regularization.

Define $h(t) := \ln(\max\{t, \xi\})$ for $t \ge 0$. We show that h is globally Lipschitz with constant $1/\xi$.

Indeed, h is differentiable on $(0, \infty) \setminus \{\xi\}$ with

$$h'(t) = \begin{cases} 0, & t < \xi, \\ \frac{1}{t}, & t > \xi. \end{cases}$$

Hence, $|h'(t)| \leq 1/\xi$ for all $t \neq \xi$. Since h is continuous and piecewise C^1, the mean value theorem on each region implies

$$|h(t) - h(s)| \leq \frac{1}{\xi}|t - s| \quad \forall t, s \geq 0.$$

Step 7: Conclusion.

Applying the Lipschitz property of h together with Step 5 yields

$$\left| \ln \Delta^*_\xi f(x) - \ln \Delta^*_\xi f(y) \right| \leq \frac{1}{\xi} \left| \|df_x\|_{\mathrm{op}} - \|df_y\|_{\mathrm{op}} \right| \leq \frac{C}{\xi} d_M(x, y).$$

This completes the proof.

This intrinsic framework establishes a coordinate-free, geometrically controlled notion of sensitivity. The operator Δ^* encodes first-order variation of the query along admissible geodesics, while bounded Hessian curvature ensures multiplicative stability of the local sensitivity profile.

4. Mechanism and privacy analysis

In this section, we establish a rigorous privacy analysis of the proposed ON–ALM. Throughout this analysis, we operate under the geometric normalization introduced in Assumption 1.

4.1. Multivariate Laplace noise and tail bounds

To calibrate the mechanism, we utilize the k-dimensional Laplace distribution. For a scale parameter $b > 0$, the multivariate Laplace density is defined as

$$p_b(z) = \frac{1}{(2b)^k} \exp\left(-\frac{\|z\|_1}{b} \right), \qquad z \in \mathbb{R}^k.$$

If a random vector $Z \sim \mathrm{Lap}_k(b)$, its coordinates are independent and identically distributed as $\mathrm{Lap}(b)$. Consequently, the ℓ_1-norm of the noise vector scales according to a Gamma distribution:

$$\frac{\|Z\|_1}{b} \sim \Gamma(k, 1).$$

This structural connection between the ℓ_1 geometry of the Laplace noise and the Gamma distribution allows us to establish the exact tail bounds required for the typical set extraction in (ε, δ)-differential privacy.

Lemma 2 (Gamma Tail Bound)**.** *Let $Z \sim \mathrm{Lap}_k(b)$ and let $\delta \in (0, 1)$. If $\Lambda_{k,\delta}$ denotes the $(1-\delta)$-quantile of the standard Gamma distribution $\Gamma(k, 1)$, then*

$$\mathbb{P}(\|Z\|_1 \leq b\, \Lambda_{k,\delta}) = 1 - \delta.$$

Proof. Since the scaled norm $\|Z\|_1/b \sim \Gamma(k, 1)$ is an absolutely continuous random variable, the exact tail probability follows directly from evaluating its cumulative distribution function at the defined $(1-\delta)$-quantile [13].

4.2. Mechanism definition

Definition 4 (ON–ALM Mechanism)**.** *Fix $\varepsilon > 0$ and $\delta \in (0, 1)$. Define the scale*

$$b(x) := \frac{e^{2\beta}\,\Delta^*_\xi f(x)}{\varepsilon - k\beta - \Lambda_{k,\delta}(e^\beta - 1)}.$$

Assume the feasibility condition

$$\varepsilon > k\beta + \Lambda_{k,\delta}(e^\beta - 1),$$

so that $b(x) > 0$.

The ON–ALM mechanism outputs

$$\mathcal{M}(x) = f(x) + Z, \qquad Z \sim \mathrm{Lap}_k(b(x)).$$

Remark 6 (Feasibility Condition: Parameter Guidance and Admissible Range of β)**.** *The feasibility condition $\varepsilon > k\beta + \Lambda_{k,\delta}(e^\beta - 1)$ imposes a joint constraint on the stability parameter β, the output dimension k, and the privacy budget ε. We provide explicit guidance for typical choices of privacy parameters.*

Admissible range for β. *Since $e^\beta - 1 > 0$ for any $\beta > 0$, and $\Lambda_{k,\delta} > k$ for $\delta < 1/2$ (as the $(1-\delta)$-quantile of $\Gamma(k, 1)$ exceeds its mean k when $\delta < 1/2$), the condition requires*

$$\varepsilon > \beta(k + \Lambda_{k,\delta}) + O(\beta^2),$$

so an explicit upper bound on β in the small-β regime is

$$\beta < \frac{\varepsilon}{k + \Lambda_{k,\delta}}.$$

Using the asymptotic expansion $\Lambda_{k,\delta} = k + O(\sqrt{k\ln(1/\delta)})$, this simplifies to

$$\beta < \frac{\varepsilon}{2k + O(\sqrt{k\ln(1/\delta)})}.$$

Concrete numerical examples.

1) Low-dimensional, strong privacy*: $k = 1$, $\varepsilon = 1$, $\delta = 10^{-5}$. The quantile $\Lambda_{1,10^{-5}} \approx 11.5$. The feasibility condition requires $\beta < \varepsilon/(1 + 11.5) \approx 0.08$. Thus, $\beta \le 0.05$ is safely admissible, corresponding to at most 5% relative variation of $\Delta^*_\xi f$ per unit geodesic distance.*
2) Moderate dimension, standard privacy*: $k = 5$, $\varepsilon = 1$, $\delta = 10^{-3}$. Here, $\Lambda_{5,10^{-3}} \approx 14.79$ ($(1-\delta)$-quantile of $\Gamma(5, 1)$). The linear bound gives $\beta < 1/(5 + 14.79) \approx 0.051$.*
3) Relaxed privacy*: $k = 5$, $\varepsilon = 3$, $\delta = 10^{-3}$. The bound gives $\beta < 3/(5 + 14.79) \approx 0.152$, allowing a substantially wider range of sensitivity variation.*

Design implication. *In practice, β should be estimated from the data or bounded analytically using the Hessian bound $\beta \le C/\xi$ (see the Geometric Interpretation subsection). If the estimated β does not satisfy the feasibility condition, one may either (i) increase ε by weakening the privacy requirement, (ii) reduce k by aggregating query components, or (iii) increase ξ to tighten the regularization floor, thereby reducing β. The linear scaling $\beta = O(\varepsilon/k)$ highlights that ON–ALM is best suited for low- to moderate-dimensional queries under standard privacy budgets ($\varepsilon \in [0.5, 3]$), which is consistent with current best practice in the differential privacy literature.*

Remark (Reduction to Global Sensitivity). If $\Delta^*_\xi f(x)$ is constant on M, i.e.,

$$\Delta^*_\xi f(x) \equiv \Delta,$$

then the sensitivity is globally uniform and exhibits no spatial variation. In this case, the stability parameter can be taken as $\beta = 0$, which implies $e^\beta - 1 = 0$. The scale reduces to

$$b(x) = \frac{\Delta}{\varepsilon},$$

recovering the classical Laplace mechanism.

For $\beta > 0$, the additional terms in the denominator and the factor $e^{2\beta}$ in the numerator quantify the deviation from global uniformity and represent the cost of enforcing stability under spatially varying sensitivity.

Lemma 3 (Privacy Loss Bound)**.** *Let $x, x' \in M$ with $d_M(x, x') \leq 1$. Let $b = b(x)$ and $b' = b(x')$. For any $y \in \mathbb{R}^k$, the privacy loss*

$$L(y) := \ln \frac{p_x(y)}{p_{x'}(y)}$$

satisfies

$$L(y) = k \ln\left(\frac{b'}{b}\right) + \frac{\|y - f(x')\|_1}{b'} - \frac{\|y - f(x)\|_1}{b}.$$

Moreover,

$$L(y) \leq k \ln\left(\frac{b'}{b}\right) + \frac{\|f(x) - f(x')\|_1}{b'} + \|y - f(x)\|_1 \left|\frac{1}{b'} - \frac{1}{b}\right|. \tag{4.1}$$

Proof. By definition, $p_x(y) = (2b)^{-k} \exp(-\|y - f(x)\|_1/b)$ and $p_{x'}(y) = (2b')^{-k} \exp(-\|y - f(x')\|_1/b')$. Taking the ratio and its logarithm directly yields the identity.

For the upper bound, apply the triangle inequality

$$\|y - f(x')\|_1 \leq \|y - f(x)\|_1 + \|f(x) - f(x')\|_1,$$

so $\|y - f(x')\|_1/b' \leq \|y - f(x)\|_1/b' + \|f(x) - f(x')\|_1/b'$. Substituting into the identity:

$$\begin{aligned} L(y) &= k \ln\tfrac{b'}{b} + \tfrac{\|y-f(x')\|_1}{b'} - \tfrac{\|y-f(x)\|_1}{b} \\ &\leq k \ln\tfrac{b'}{b} + \tfrac{\|f(x)-f(x')\|_1}{b'} + \|y - f(x)\|_1\left(\tfrac{1}{b'} - \tfrac{1}{b}\right). \end{aligned}$$

Since the last factor $(1/b' - 1/b)$ may be negative (when $b' > b$) or positive (when $b' < b$), we bound it from above by its absolute value $|1/b' - 1/b|$, obtaining (4.1).

4.3. Main privacy theorem

Theorem 1 ((ε, δ)-Differential Privacy)**.** *Let $\mathcal{M}$ be the ON–ALM mechanism. Under Assumption 1, and assuming the feasibility condition*

$$\varepsilon > k\beta + \Lambda_{k,\delta}(e^\beta - 1), \tag{4.2}$$

$\mathcal{M}$ satisfies (ε, δ)-differential privacy.

Proof. Let $D, D' \in \mathcal{D}$ be adjacent datasets and set $x = \Phi(D)$, $x' = \Phi(D') \in M$ for their respective embeddings. By Assumption 1(ii),(iii), $d_M(x, x') \le 1$. Denote $b = b(x)$ and $b' = b(x')$ for their respective scale parameters.

Symmetry and measurability. We must show that for every Borel measurable set $S \subset \mathbb{R}^k$, both inequalities

$$\mathbb{P}(\mathcal{M}(x) \in S) \le e^{\varepsilon}\mathbb{P}(\mathcal{M}(x') \in S) + \delta \quad \text{and} \quad \mathbb{P}(\mathcal{M}(x') \in S) \le e^{\varepsilon}\mathbb{P}(\mathcal{M}(x) \in S) + \delta$$

hold. We prove the first direction explicitly and then establish the second direction, showing that the argument is fully symmetric.

Typical set and tail bound. Define the typical set around $f(x)$ as

$$\mathcal{T}_x = \{y \in \mathbb{R}^k : \|y - f(x)\|_1 \le b\,\Lambda_{k,\delta}\},$$

where $b = b(x) > 0$ is determined by x. For each fixed x, $\mathcal{T}_x$ is a closed ℓ_1-ball of radius $b\Lambda_{k,\delta}$ centered at $f(x) \in \mathbb{R}^k$; as a closed ball in $\mathbb{R}^k$, it is a Borel set. Let $Z \sim \mathrm{Lap}_k(b)$ so that $\mathcal{M}(x) = f(x) + Z$. By Lemma 2, $\mathbb{P}(\|Z\|_1 \le b\,\Lambda_{k,\delta}) = 1 - \delta$, so

$$\mathbb{P}(\mathcal{M}(x) \in \mathcal{T}_x) = 1 - \delta, \qquad \text{in particular,} \quad \mathbb{P}(\mathcal{M}(x) \notin \mathcal{T}_x) = \delta.$$

Justification of the Gamma tail. To see why $\|Z\|_1/b \sim \Gamma(k, 1)$: since $Z_i \overset{\text{iid}}{\sim} \mathrm{Lap}(b)$, each $|Z_i|/b \sim \mathrm{Exp}(1)$, and the sum of k independent Exp(1) random variables is $\Gamma(k, 1)$ by the standard convolution formula for exponential random variables. Lemma 2 formalizes this and provides the exact quantile $\Lambda_{k,\delta}$.

Scale and sensitivity bounds. The scale $b(x)$ is directly proportional to the regularized sensitivity $\Delta^*_\xi f(x)$. Since they differ by a positive multiplicative constant, their logarithms differ by an additive constant, which cancels in differences. Therefore, the β-log-Lipschitz property of $\Delta^*_\xi f$ implies that b is also β-log-Lipschitz, yielding

$$e^{-\beta} \le \frac{b'}{b} \le e^{\beta}.$$

In particular, we have $\ln(b'/b) \le \beta$, $1/b' \le e^{\beta}/b$, and $|1/b' - 1/b| \le (e^{\beta} - 1)/b$.
Derivation of scale ratio bounds. By Definition 4, $b(x) = C \cdot \Delta^*_\xi f(x)$ where

$$C = \frac{e^{2\beta}}{\varepsilon - k\beta - \Lambda_{k,\delta}(e^{\beta} - 1)} > 0$$

is a universal constant depending only on $(\varepsilon, \beta, k, \delta)$, not on x. Hence, $b(x') = C \cdot \Delta^*_\xi f(x')$ with the *same* C. Then,

$$\ln\frac{b'}{b} = \ln \Delta^*_\xi f(x') - \ln \Delta^*_\xi f(x),$$

and the β-log-Lipschitz condition with $d_M(x, x') \le 1$ gives $|\ln(b'/b)| \le \beta$, so $e^{-\beta} \le b'/b \le e^{\beta}$. The bound $1/b' \le e^{\beta}/b$ follows directly from $b' \ge e^{-\beta}b$. For the absolute-value bound $|1/b' - 1/b| \le (e^{\beta} - 1)/b$, consider two cases. *Case 1* ($b' \le b$): then $1/b' \ge 1/b > 0$, so $1/b' - 1/b = (b - b')/(bb') > 0$ and $|1/b' - 1/b| = 1/b' - 1/b$. Since $1/b' \le e^{\beta}/b$ (from $b' \ge e^{-\beta}b$), we get $|1/b' - 1/b| = 1/b' - 1/b \le e^{\beta}/b - 1/b = (e^{\beta} - 1)/b$. *Case 2* ($b' > b$): then $1/b' - 1/b < 0$, so $|1/b' - 1/b| = 1/b - 1/b' \le 1/b - 1/(e^{\beta}b) = (e^{\beta} - 1)/(e^{\beta}b) \le (e^{\beta} - 1)/b$. In both cases, $|1/b' - 1/b| \le (e^{\beta} - 1)/b$.

To bound the deterministic variation of the query, let $\gamma : [0, \ell] \to M$ be a minimizing unit-speed geodesic connecting x and x', where $\ell = d_M(x, x') \le 1$. By the injectivity radius assumption, such a minimizing geodesic exists. By the intrinsic sensitivity bound (Proposition 3) and using that $\Delta^* f \le \Delta^*_\xi f$, we have

$$\|f(x) - f(x')\|_1 \le \int_0^\ell \Delta^* f(\gamma(t))\, dt \le \ell \sup_{t\in[0,\ell]} \Delta^*_\xi f(\gamma(t)).$$

Applying the β-log-Lipschitz property to the pair $(x, \gamma(t))$: since γ is a unit-speed *minimizing* geodesic and $t \le \ell < \mathrm{inj}(x)$, we have $d_M(x, \gamma(t)) = t$ (the arc length equals the geodesic distance). Hence

$$\ln \Delta^*_\xi f(\gamma(t)) - \ln \Delta^*_\xi f(x) \le \beta\, d_M(x, \gamma(t)) = \beta t,$$

giving $\sup_{t\in[0,\ell]} \Delta^*_\xi f(\gamma(t)) \le e^{\beta\ell} \Delta^*_\xi f(x) \le e^{\beta} \Delta^*_\xi f(x)$ (using $\ell \le 1$). We conclude

$$\|f(x) - f(x')\|_1 \le e^{\beta} \Delta^*_\xi f(x).$$

Why $\sup_t \Delta^*_\xi f(\gamma(t)) \le e^{\beta} \Delta^*_\xi f(x)$. The log-Lipschitz condition states that for all $t \in [0, \ell]$,

$$\ln \Delta^*_\xi f(\gamma(t)) \le \ln \Delta^*_\xi f(x) + \beta\, d_M(\gamma(t), x) = \ln \Delta^*_\xi f(x) + \beta t,$$

since γ is a unit-speed geodesic and $d_M(\gamma(t), x) = t$. Exponentiating yields $\Delta^*_\xi f(\gamma(t)) \le e^{\beta t} \Delta^*_\xi f(x) \le e^{\beta\ell} \Delta^*_\xi f(x) \le e^{\beta} \Delta^*_\xi f(x)$ for all $t \in [0, \ell]$, using $\ell \le 1$.

Privacy loss bound. We now evaluate the privacy loss $L(y)$ for any fixed output $y \in \mathcal{T}_x$. Using the privacy loss decomposition from Lemma 3, we obtain

$$L(y) \le k \ln\!\left(\frac{b'}{b}\right) + \frac{\|f(x) - f(x')\|_1}{b'} + \|y - f(x)\|_1 \left|\frac{1}{b'} - \frac{1}{b}\right|.$$

Substituting (i) $k \ln(b'/b) \le k\beta$, (ii) $\|y - f(x)\|_1 \le b\Lambda_{k,\delta}$ with $|1/b' - 1/b| \le (e^{\beta} - 1)/b$, and (iii) the query variation bound $\|f(x) - f(x')\|_1 \le e^{\beta} \Delta^*_\xi f(x)$ together with $1/b' \le e^{\beta}/b$:

$$L(y) \;\le\; k\beta \;+\; \underbrace{\frac{\|f(x) - f(x')\|_1}{b'}}_{\le e^{\beta}\Delta^*_\xi f(x)/b' \,\le\, e^{2\beta}\Delta^*_\xi f(x)/b} \;+\; \underbrace{b\,\Lambda_{k,\delta} \cdot \frac{e^{\beta} - 1}{b}}_{=\Lambda_{k,\delta}(e^{\beta}-1)} \;\le\; k\beta + \frac{e^{2\beta}\Delta^*_\xi f(x)}{b} + \Lambda_{k,\delta}(e^{\beta} - 1).$$

By Definition 4, the scale is chosen so that

$$\frac{e^{2\beta}\Delta^*_\xi f(x)}{b(x)} = \varepsilon - k\beta - \Lambda_{k,\delta}(e^{\beta} - 1).$$

Substituting this identity into the previous bound perfectly cancels the terms involving $k\beta$ and $\Lambda_{k,\delta}(e^{\beta} - 1)$, giving

$$L(y) \le \varepsilon \qquad \text{for all } y \in \mathcal{T}_x.$$

Since the Laplace density is strictly positive everywhere on $\mathbb{R}^k$, we have $p_{x'}(y) > 0$ for all y, so the inequality $L(y) \le \varepsilon$ is equivalent to $p_x(y) \le e^{\varepsilon} p_{x'}(y)$. Hence $p_x(y) \le e^{\varepsilon} p_{x'}(y)$ for all $y \in \mathcal{T}_x$.

Explicit cancellation. To see the cancellation in full detail, substituting $e^{2\beta}\Delta^*_\xi f(x)/b = \varepsilon - k\beta - \Lambda_{k,\delta}(e^\beta - 1)$ into

$$L(y) \le k\beta + \frac{e^{2\beta}\Delta^*_\xi f(x)}{b} + \Lambda_{k,\delta}(e^\beta - 1)$$

gives $L(y) \le k\beta + [\varepsilon - k\beta - \Lambda_{k,\delta}(e^\beta - 1)] + \Lambda_{k,\delta}(e^\beta - 1) = \varepsilon$, as claimed. The scale formula in Definition 4 is precisely reverse-engineered so that this cancellation occurs.

Final step (Measure-theoretic decomposition). Let $S \subset \mathbb{R}^k$ be any Borel measurable set. We decompose the probability measure:

$$\begin{aligned}
\mathbb{P}(\mathcal{M}(x) \in S) &= \mathbb{P}(\mathcal{M}(x) \in S \cap \mathcal{T}_x) + \mathbb{P}(\mathcal{M}(x) \in S \setminus \mathcal{T}_x) \\
&= \int_{S \cap \mathcal{T}_x} p_x(y)\, dy + \mathbb{P}(\mathcal{M}(x) \notin \mathcal{T}_x) \\
&\le \int_{S \cap \mathcal{T}_x} e^\varepsilon p_{x'}(y)\, dy + \mathbb{P}(\mathcal{M}(x) \notin \mathcal{T}_x) \\
&\le e^\varepsilon \mathbb{P}(\mathcal{M}(x') \in S) + \delta,
\end{aligned}$$

where in the last step we used $\int_{S \cap \mathcal{T}_x} e^\varepsilon p_{x'}(y)\, dy \le e^\varepsilon \int_S p_{x'}(y)\, dy = e^\varepsilon\, \mathbb{P}(\mathcal{M}(x') \in S)$ (valid since $p_{x'} \ge 0$), and $\mathbb{P}(\mathcal{M}(x) \notin \mathcal{T}_x) = \delta$ (by Lemma 2).

Reverse inequality. It remains to show $\mathbb{P}(\mathcal{M}(x') \in S) \le e^\varepsilon\, \mathbb{P}(\mathcal{M}(x) \in S) + \delta$. Define the typical set centered at x':

$$\mathcal{T}_{x'} = \{y \in \mathbb{R}^k : \|y - f(x')\|_1 \le b' \Lambda_{k,\delta}\},$$

with $b' = b(x')$. By Lemma 2, $\mathbb{P}(\mathcal{M}(x') \notin \mathcal{T}_{x'}) = \delta$.

For $y \in \mathcal{T}_{x'}$ we bound the *reverse* privacy loss. From the lemma identity with roles of x and x' interchanged:

$$L'(y) := \ln \frac{p_{x'}(y)}{p_x(y)} = k \ln \frac{b}{b'} + \frac{\|y - f(x)\|_1}{b} - \frac{\|y - f(x')\|_1}{b'}.$$

Applying the triangle inequality to $\|y - f(x)\|_1 \le \|y - f(x')\|_1 + \|f(x') - f(x)\|_1$ gives

$$L'(y) \le k \ln \frac{b}{b'} + \frac{\|f(x') - f(x)\|_1}{b} + \|y - f(x')\|_1 \left(\frac{1}{b} - \frac{1}{b'}\right).$$

We bound each term for $y \in \mathcal{T}_{x'}$:

- *Scale ratio:* $k \ln(b/b') \le k\beta$ (since $|\ln(b'/b)| \le \beta$).

- *Query variation:* Define $\sigma(t) = \gamma(\ell - t)$; then σ is a unit-speed minimizing geodesic from x' to x with $d_M(x', \sigma(t)) = t$. Applying Proposition 3 to σ, we get $\|f(x') - f(x)\|_1 \leq \int_0^\ell \Delta_\xi^* f(\sigma(t))\,dt$. The β-log-Lipschitz bound along σ gives $\Delta_\xi^* f(\sigma(t)) \leq e^{\beta t}\Delta_\xi^* f(x')$, so $\sup_t \Delta_\xi^* f(\sigma(t)) \leq e^{\beta\ell}\Delta_\xi^* f(x') \leq e^{\beta}\Delta_\xi^* f(x')$. Hence:

$$\frac{\|f(x') - f(x)\|_1}{b} \leq \frac{e^{\beta}\,\Delta_\xi^* f(x')}{b} \leq \frac{e^{2\beta}\,\Delta_\xi^* f(x')}{b'},$$

  where the last step uses $1/b \leq e^{\beta}/b'$ (from $b' \leq e^{\beta} b$).
- *Last term:* For $y \in \mathcal{T}_{x'}$, $\|y - f(x')\|_1 \leq b'\Lambda_{k,\delta}$. The reversed coefficient satisfies $|1/b - 1/b'| \leq (e^{\beta} - 1)/b'$: in Case 1 ($b \geq b'$): $1/b \leq e^{\beta}/b'$ (from $b' \geq e^{-\beta} b$), so $1/b - 1/b' \leq e^{\beta}/b' - 1/b' = (e^{\beta} - 1)/b'$; in Case 2 ($b < b'$): $1/b - 1/b' < 0 \leq (e^{\beta} - 1)/b'$. Hence:

$$\|y - f(x')\|_1 \left|\frac{1}{b} - \frac{1}{b'}\right| \leq b'\Lambda_{k,\delta} \cdot \frac{e^{\beta} - 1}{b'} = \Lambda_{k,\delta}(e^{\beta} - 1).$$

- *Cancellation:* By Definition 4, $e^{2\beta}\Delta_\xi^* f(x')/b(x') = \varepsilon - k\beta - \Lambda_{k,\delta}(e^{\beta} - 1)$, so

$$L'(y) \leq k\beta + [\varepsilon - k\beta - \Lambda_{k,\delta}(e^{\beta} - 1)] + \Lambda_{k,\delta}(e^{\beta} - 1) = \varepsilon.$$

Since $p_x(y) > 0$ for all y, $L'(y) \leq \varepsilon$ implies $p_{x'}(y) \leq e^{\varepsilon} p_x(y)$ for all $y \in \mathcal{T}_{x'}$. Applying the measure-theoretic decomposition with $\mathcal{T}_{x'}$ in place of $\mathcal{T}_x$ and $\mathbb{P}(\mathcal{M}(x') \notin \mathcal{T}_{x'}) = \delta$ yields

$$\mathbb{P}(\mathcal{M}(x') \in S) \leq e^{\varepsilon}\,\mathbb{P}(\mathcal{M}(x) \in S) + \delta,$$

completing the proof of (ε, δ)-differential privacy.

Correctness of the final inequality. The key step is $\int_{S \cap \mathcal{T}_x} p_x(y)\,dy \leq e^{\varepsilon} \int_{S \cap \mathcal{T}_x} p_{x'}(y)\,dy$. This follows because $p_x(y) \leq e^{\varepsilon} p_{x'}(y)$ holds pointwise for every $y \in \mathcal{T}_x$, as established above. Extending the integration from $S \cap \mathcal{T}_x$ to S is valid because all integrands are nonnegative, so $\int_{S \cap \mathcal{T}_x} e^{\varepsilon} p_{x'}(y)\,dy \leq e^{\varepsilon} \int_S p_{x'}(y)\,dy = e^{\varepsilon}\mathbb{P}(\mathcal{M}(x') \in S)$.

We summarize the proposed ON–ALM mechanism in Algorithm 1. The procedure highlights that, despite the rigorous geometric formulation, the mechanism admits a closed-form and highly implementable calibration rule.

Algorithm 1 Operator normed adaptive Laplace mechanism (ON–ALM)

Input: Dataset $D \in \mathcal{D}$, query $f : M \to \mathbb{R}^k$, privacy parameters (ε, δ), regularization parameters $\beta, \xi > 0$.
Output: Private estimate $\tilde{f} \in \mathbb{R}^k$.

1: Compute the embedded point $x \leftarrow \Phi(D) \in M$.
2: Compute (or estimate) the differential df_x.
3: Set
$$\Delta^*_\xi f(x) = \max\{\|df_x\|_{\mathrm{op}}, \xi\}.$$
4: Compute the $(1-\delta)$-quantile $\Lambda_{k,\delta}$ of the $\Gamma(k, 1)$ distribution.
5: Set
$$b = \frac{e^{2\beta}\Delta^*_\xi f(x)}{\varepsilon - k\beta - \Lambda_{k,\delta}(e^\beta - 1)}.$$
6: Sample $Z \sim \mathrm{Lap}_k(b)$ with density $p(z) \propto e^{-\|z\|_1/b}$.
7: Set
$$\tilde{f} = f(x) + Z.$$
8: **return** $\tilde{f}$

Remark 7 (Computational Tractability and Dataset-Size Independence). *Calibrating noise in the classical smooth sensitivity framework [2] requires computing a global supremum over the data manifold,*
$$S(x) = \sup_{y \in M} \Delta(y) e^{-\beta d_M(x,y)}.$$
For high-dimensional or complex datasets, solving this global optimization problem is computationally prohibitive.

The ON–ALM framework bypasses this global search. By Proposition 4, a bounded Riemannian Hessian implies that the regularized local sensitivity $\Delta^*_\xi f(x)$ is intrinsically β-log-Lipschitz. Thus, Algorithm 1 relies solely on the local *evaluation of the operator norm $\|df_x\|_{\mathrm{op}}$ at the given point x.*

This geometric reduction shifts the calibration step from a global optimization task to a local differential computation. Because the complexity of evaluating $\|df_x\|_{\mathrm{op}}$ is independent of the dataset size, the mechanism provides a scalable approach to instance-adaptive privacy.

Remark 8 (Implementation Strategies and Complexity of the Operator Norm). *While Algorithm 1 avoids a global optimization, the local computation of $\|df_x\|_{\mathrm{op}} = \sup_{\|v\|_g=1} \|df_x(v)\|_1$ can be nontrivial for general nonlinear queries. We characterize the cost and provide concrete implementation strategies.*

Closed-form cases. *For many practically relevant queries, the operator norm reduces to a closed-form expression:*

- Linear queries $f(x) = Wx$ *(where $W \in \mathbb{R}^{k\times m}$): The differential is constant, $df_x = W$, and the $(\ell_2 \to \ell_1)$ operator norm is*
$$\|W\|_{2\to1} = \max_{\|v\|_2=1} \|Wv\|_1 = \max_{\|v\|_2=1} \sum_{i=1}^{k} |W_{i,:}v|.$$

This is not *simply* $\max_j \|W_{j,:}\|_2$*; the correct value can be up to* $\sqrt{k}$ *times larger (e.g., for* $W = I_k$*,* $\|W\|_{2\to1} = \sqrt{k}$ *while* $\max_j \|W_{j,:}\|_2 = 1$*). In general,* $\|W\|_{2\to1}$ *can be computed exactly as the maximum* ℓ_1 *norm of* Wv *over the unit sphere, which equals the maximum column norm of* W^T *in the* ℓ_∞*-to-*ℓ_1 *sense. For diagonal* $W = \mathrm{diag}(w_1, \ldots, w_k)$*, it reduces to* $\|w\|_2$ *(Cauchy-Schwarz achievable by* $v = w/\|w\|_2$*), computable in* $O(k)$*. For general* W*, a tight upper bound* $\|W\|_{2\to1} \le \|W\|_F$ *(Frobenius norm) is computable in* $O(km)$*.*

- Scalar quadratic queries $f(x) = x^\top Ax$*:* $\|df_x\|_{\mathrm{op}} = \|Ax\|_2$*, requiring one matrix-vector product in* $O(m^2)$.
- Neural network queries with smooth activations*: The differential* df_x *is the Jacobian matrix* $J_f(x) \in \mathbb{R}^{k\times m}$*. The mixed* $(\ell_2 \to \ell_1)$ *operator norm equals* $\max_{\|v\|_2=1} \|J_f(x)v\|_1 \le \sum_{i=1}^k \|J_f(x)_{i,:}\|_2$*. The Jacobian can be computed via automatic differentiation in* $O(k)$ *backward passes, each of cost* $O(P)$ *where* P *is the number of parameters.*

Numerical approximation via power iteration. *For general differentiable queries,* $\|df_x\|_{\mathrm{op}}$ *can be approximated using the following iterative procedure. Note that since* ℓ_1 *and* ℓ_2 *norms are equivalent on* $\mathbb{R}^k$ *up to dimension-dependent constants (*$\|v\|_1 \le \sqrt{k}\|v\|_2$*), one may instead compute the spectral norm* $\|J_f(x)\|_2$ *and apply the inequality:*

$$\|J_f(x)\|_2 \le \|df_x\|_{\mathrm{op}} \le \sqrt{k}\,\|J_f(x)\|_2.$$

The spectral norm $\|J_f(x)\|_2$ *equals the largest singular value of the Jacobian and can be estimated via a randomized power iteration in* $O(T \cdot \mathrm{JVP})$ *time, where* T *is the number of iterations (typically* $T = 10$–20*) and* JVP *is the cost of a Jacobian-vector product via automatic differentiation. This yields a privacy-valid upper bound when the* $\sqrt{k}$ *factor is absorbed into the noise scale.*

Complexity summary. *Let* $m = \dim(M)$ *and* k *be the output dimension. The dominant costs per dataset point are:*

1) Embedding evaluation $\Phi(D)$*:* $O(Nm)$ *for tabular mean embeddings.*

2) Jacobian computation: $O(k \cdot \text{backward\_pass})$ *via automatic differentiation.*

3) Operator norm estimation: $O(T \cdot k \cdot m)$ *via power iteration, with* $T \ll \min(k, m)$*.*

4) Noise sampling: $O(k)$*.*

Total per-query cost is therefore $O(k \cdot m + T \cdot k \cdot m) = O(Tkm)$*, which is independent of dataset size* N *and linear in both manifold and output dimension. This is substantially cheaper than smooth sensitivity computation, which requires* $O(N)$ *evaluations over neighboring datasets in its discrete approximation.*

SPD manifold case. *When* $M = \mathcal{S}^d_{++}$ *and* $f(P) = \mathrm{tr}(WP)$ *for a fixed symmetric matrix* W*, the Riemannian differential under the affine-invariant metric* $g_P(U, V) = \mathrm{tr}(P^{-1}UP^{-1}V)$ *satisfies* $df_P(V) = \mathrm{tr}(WV)$*. The operator norm with respect to the Riemannian tangent-space norm* $\|V\|_P = \|P^{-1/2}VP^{-1/2}\|_F$ *is*

$$\|df_P\|_{\mathrm{op}} = \sup_{\|P^{-1/2}VP^{-1/2}\|_F=1} |\mathrm{tr}(WV)| = \|P^{1/2}WP^{1/2}\|_F,$$

where the last equality follows by substituting $U = P^{-1/2}VP^{-1/2}$ *and applying Cauchy-Schwarz on* $\langle P^{1/2}WP^{1/2}, U\rangle_F$*. Note that this is* $\|P^{1/2}WP^{1/2}\|_F$*,* not $\|WP^{1/2}\|_F$*; these coincide only when* W *and* P *commute (e.g., both diagonal). The correct formula is computable in* $O(d^3)$ *via the symmetric matrix square root.*

4.4. Examples

Example 3 (Anisotropic Queries on a Flat Manifold)**.** *To build intuition before applying our framework to highly curved spaces, consider $M \subset \mathbb{R}^n$ as the open unit ball equipped with the induced Euclidean metric (which is a flat Riemannian manifold), and consider scalar-valued queries ($k = 1$).*

Linear query. *If $f(x) = \langle w, x\rangle$, then the differential corresponds to the standard gradient. The operator norm from $(T_xM, \|\cdot\|_2)$ to $(\mathbb{R}, |\cdot|)$ yields the dual norm*

$$\Delta^* f(x) = \|w\|_2.$$

Since the sensitivity is constant across M, ON–ALM seamlessly reduces to the classical global Laplace mechanism.

Quadratic query. *Let $A \in \mathbb{R}^{n\times n}$ be a symmetric matrix. If $g(x) = \frac{1}{2}x^\top Ax$, then*

$$\nabla g(x) = Ax, \qquad \Delta^* g(x) = \|Ax\|_2.$$

In this case, the sensitivity is strictly position-dependent. ON–ALM adapts to the local geometry by injecting significantly less noise in regions where $\|Ax\|_2$ is small, providing higher utility than a global worst-case bound (which would pessimistically depend on the maximum singular value of A). As illustrated in Figure 1, ON–ALM seamlessly adapts to this directional variability.

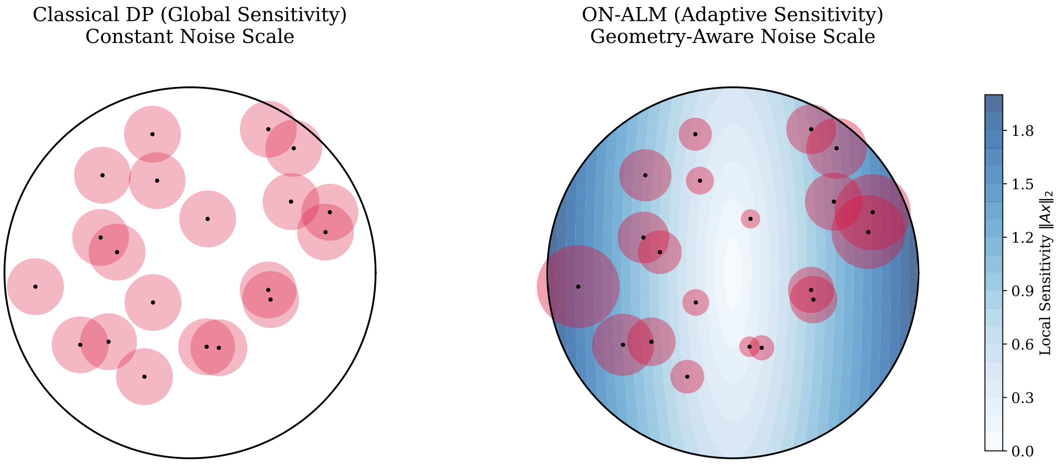

Figure 1. Visualizing instance-adaptivity on the unit disk ($n = 2$). Left: Linear query with constant sensitivity; the required Laplace noise scale (red circles) is uniform across the domain. Right: Quadratic query $g(x) = \frac{1}{2}x^\top Ax$ with an anisotropic matrix A. The background contour maps the local sensitivity $\|Ax\|_2$. ON–ALM dynamically calibrates the noise scale to the local geometry, injecting significantly less noise near the origin.

4.5. Dimensionality and stability considerations

The feasibility condition

$$\varepsilon > k\beta + \Lambda_{k,\delta}(e^\beta - 1)$$

explicitly connects the output dimension k, the stability parameter β, and the privacy level ε.

Recall that $\Lambda_{k,\delta}$ is the $(1-\delta)$-quantile of the $\Gamma(k,1)$ distribution. For typical privacy applications ($\delta \ll 0.5$), this quantile lies strictly in the right tail. Standard concentration results for gamma random variables (e.g., Laurent-Massart bounds) yield

$$\Lambda_{k,\delta} = k + O\Big(\sqrt{k\ln(1/\delta)}\Big) \qquad \text{as } k \to \infty.$$

To maintain feasibility, the mechanism operates in the small-β regime ($\beta \to 0$), where the expansion $e^{\beta} - 1 \approx \beta$ holds. Substituting the asymptotic expansion of the quantile into the feasibility condition gives

$$k\beta + \Lambda_{k,\delta}(e^{\beta} - 1) = k\beta + k(e^{\beta} - 1) + O\Big((e^{\beta} - 1)\sqrt{k\ln(1/\delta)}\Big).$$

For fixed δ and small β, the residual error term is dominated by the linear growth in k. Thus, the privacy constraint behaves asymptotically as

$$\varepsilon \gtrsim k(\beta + e^{\beta} - 1).$$

Using the approximation $e^{\beta} - 1 \approx \beta$, this threshold simplifies to

$$\varepsilon \gtrsim 2k\beta.$$

This scaling highlights a structural limit: for high-dimensional outputs, instance-adaptive mechanisms must compensate for both the dimensional growth in the Laplace tail and the spatial variability of the sensitivity encoded by β. Consequently, ON–ALM is well-suited for low- or moderate-dimensional geometric queries, or settings where β is small. This linear-in-dimension behavior is consistent with known dimensional dependencies in (ε,δ)-differential privacy for ℓ_1-calibrated mechanisms.

4.5.1. Numerical sensitivity analysis

To quantify the feasible operating regime of ON–ALM, Table 1 reports the maximum admissible β across a range of output dimensions k, privacy budgets ε, and failure probability $\delta = 10^{-5}$. Values in the table are computed from the exact feasibility condition using standard $\Gamma(k,1)$ quantiles.

Table 1. Maximum admissible stability parameter $\beta_{\max}$ satisfying the feasibility condition $\varepsilon > k\beta + \Lambda_{k,\delta}(e^{\beta} - 1)$, for $\delta = 10^{-5}$. Values are computed exactly by bisection on the closed-form condition. Bold entries indicate settings where $\beta_{\max} \geq 0.05$, regarded as practically useful. The quantile values are $\Lambda_1 = 11.51$, $\Lambda_2 = 14.24$, $\Lambda_5 = 20.65$, $\Lambda_{10} = 29.52$, $\Lambda_{20} = 45.04$ (with $\delta = 10^{-5}$).

	Output dimension k				
ε	$k = 1$	$k = 2$	$k = 5$	$k = 10$	$k = 20$
0.5	0.039	0.030	0.019	0.013	0.008
1.0	**0.077**	**0.060**	0.038	0.025	0.015
2.0	**0.149**	**0.117**	**0.076**	**0.050**	0.030
3.0	**0.217**	**0.171**	**0.112**	**0.074**	0.045

The table reveals a clear pattern: for scalar or low-dimensional outputs ($k \leq 5$) and moderate privacy budgets ($\varepsilon \geq 1$), ON–ALM admits meaningful sensitivity variation ($\beta_{\max} \geq 0.02$–0.06), sufficient for most smooth queries on compact manifolds. For $k \geq 10$, the feasible β range becomes very small ($\beta_{\max} < 0.02$), effectively limiting the mechanism to nearly globally uniform sensitivity.

Tension with high-dimensional motivation. The Introduction positions ON–ALM as relevant to complex representations such as deep learning embeddings and natural language models. These typically operate at high output dimensions ($k \gg 10$). In these regimes, the feasibility constraint becomes restrictive. However, we emphasize that ON–ALM is primarily designed for *query outputs* rather than intermediate representations. In practice, final statistical queries of interest (e.g., mean embedding norms, covariance traces, principal angles) are often low-dimensional ($k \leq 10$), even when the underlying data lives in high-dimensional ambient space.

Remedies for high-dimensional settings. We discuss three approaches for extending the effective operating range of ON–ALM to higher output dimensions.

1) *Dimensionality reduction via random projection.* Apply a Johnson-Lindenstrauss projection $\Pi \in \mathbb{R}^{r\times k}$ ($r \ll k$) to the query output before privatization. For a scaled Gaussian matrix $\Pi = r^{-1/2}G$ with $G_{ij} \overset{\text{iid}}{\sim} \mathcal{N}(0,1)$, the spectral norm satisfies $\|\Pi\|_{\text{op}} \leq (1+t)\sqrt{k/r}$ with probability at least $1-2e^{-rt^2/2}$ (by a standard random matrix concentration argument). Taking $t = \sqrt{2\ln(2/\eta)/r}$ yields a deterministic bound holding with probability $1-\eta$. Under this event:

$$\|\Delta^*(\Pi \circ f)(x)\|_{2\to 1} \leq \|\Pi\|_{\text{op}} \cdot \|\Delta^* f(x)\|_{2\to 1} = \|\Pi\|_{\text{op}} \cdot \Delta^* f(x),$$

where the factorization holds because $d(\Pi \circ f)_x = \Pi \circ df_x$, so $\|\Pi \circ df_x\|_{2\to 1} \leq \|\Pi\|_{\text{op}} \cdot \|df_x\|_{2\to 1}$. Privacy is preserved with the reduced output dimension r and sensitivity inflated by $\|\Pi\|_{\text{op}}$. The representation error is bounded by the Johnson–Lindenstrauss (JL) distortion, which satisfies $\|\Pi f(x) - f(x)\|^2 \leq (2\ln(2k/\eta)/r)\|f(x)\|^2$ for an appropriate projection [14].

2) *Component-wise privatization with composition.* For separable query outputs $f = (f_1, \ldots, f_k)$ with k_0-dimensional sub-blocks, apply ON–ALM independently to each block of size $k_0 \leq k_{\max}$ and compose using advanced composition theorems. Under (ε', δ')-DP per block and $\lceil k/k_0 \rceil$ blocks, the composed mechanism satisfies (ε, δ)-DP with $\varepsilon = O(\varepsilon' \sqrt{\lceil k/k_0 \rceil \ln(1/\delta)})$ by advanced composition [15].

3) *Alternative noise geometries.* For high-dimensional outputs, the ℓ_2-calibrated multivariate Gaussian mechanism may be preferable: it replaces the ℓ_1 Laplace tail with ℓ_2 Gaussian tails, where dimensional growth is $O(\sqrt{k})$ rather than $O(k)$. The geometric sensitivity operator $\Delta^* f$ defined in this work translates directly to the Gaussian setting by replacing the ℓ_1 operator norm with the spectral (Frobenius) norm of the Jacobian. Combining our Riemannian sensitivity framework with Gaussian-DP calibration [4] is an immediate direction for future work.

4.6. Geometric interpretation of the stability parameter β

The stability constant β has a direct geometric interpretation based on the second-order regularity of the query and the curvature of the underlying manifold.

Recall that if $f \in C^2(M, \mathbb{R}^k)$ satisfies

$$\|\nabla^2 f(x)\|_{\text{op}} \leq C \qquad \text{for all } x \in M,$$

and $\Delta_\xi^* f(x) \geq \xi > 0$, then $\Delta_\xi^* f$ is β-log-Lipschitz with

$$\beta \leq \frac{C}{\xi}.$$

Thus, β is bounded by the ratio of the maximal operator norm of the Riemannian Hessian to the minimal admissible sensitivity. In geometric terms:

- The Hessian bound C quantifies the second-order variation (acceleration) of the query along geodesics.
- The denominator ξ prevents the degeneration of the scale, ensuring that logarithmic derivatives remain bounded.
- If the sectional curvature of (M, g) is uniformly bounded, geodesic deviation is controlled. Via the Jacobi equations, this stabilizes the evolution of ∇f along minimizing geodesics.

Accordingly, β acts as a curvature-sensitive nonlinearity parameter. A small β corresponds to queries that are nearly affine along geodesics, while a large β indicates strong second-order variation relative to the local sensitivity magnitude.

This perspective shows that the stability parameter is an intrinsic geometric property defining how the query interacts with the Riemannian structure, rather than an abstract analytic artifact.

5. Comparison with related mechanisms

Recent developments in DP have increasingly focused on geometry-aware and manifold-based approaches. Examples include differentially private Riemannian optimization [6], privacy-preserving manifold learning and crowd sensing [7], hybrid privacy mechanisms based on Fréchet means [9], and recent manifold-oriented privacy frameworks [8, 10]. While these studies demonstrate the growing importance of geometric structures in privacy-preserving data analysis, they address substantially different objectives from the present work. The purpose of this section is therefore to position ON–ALM relative to the most closely related privacy mechanisms in terms of sensitivity adaptation, noise calibration, privacy guarantees, and geometric awareness.

We situate ON–ALM within the landscape of existing DP mechanisms, focusing on five dimensions: (i) sensitivity type used for noise calibration, (ii) noise distribution and tail behavior, (iii) DP guarantee type, (iv) whether intrinsic Riemannian geometry is exploited, and (v) computational cost. Table 2 summarizes these comparisons.

5.1. When ON–ALM outperforms competing mechanisms

vs. Classical Laplace. ON–ALM strictly dominates the classical Laplace mechanism whenever the local sensitivity $\Delta^* f(x)$ is significantly smaller than the global sensitivity $\sup_x \Delta^* f(x)$. Formally, the noise reduction factor at point x is

$$\frac{b_{\text{ON-ALM}}(x)}{b_{\text{Lap}}} = \frac{e^{2\beta}\Delta_\xi^* f(x)/(\varepsilon - k\beta - \Lambda_{k,\delta}(e^\beta - 1))}{\sup_x \Delta_\xi^* f(x)/\varepsilon} < 1$$

whenever $\Delta_\xi^* f(x) \ll \sup_x \Delta_\xi^* f(x)$ and β is small. In the quadratic query example ($f(x) = \frac{1}{2}x^\top Ax$), this ratio is $\|Ax\|_2/\|A\|_{\text{op}}$, which approaches zero near the origin.

Table 2. Comparison of ON–ALM with related differential privacy mechanisms. GS = global sensitivity; LS = local sensitivity; SS = smooth sensitivity. "Light-tailed" means finite noise variance; "Heavy-tailed" denotes Cauchy or Student-t distributions with infinite variance. ✓ = supported; × = not supported; ∼ = partially supported. The SS+Gaussian row refers to the (ε, δ)-variant of smooth sensitivity using Gaussian noise calibrated to the smooth sensitivity upper bound.

Mechanism	Sensitivity	Noise tail	DP type	Riem. geom.	Instance-adaptive
Laplace [1]	GS	Light	$(\varepsilon, 0)$	×	×
SS+Cauchy [2]	SS (global)	Heavy	$(\varepsilon, 0)$	×	✓
SS+Gaussian [2]	SS (global)	Light	(ε, δ)	×	✓
K-Norm Gradient [5]	GS (geom.)	Light	$(\varepsilon, 0)$	∼	×
Gaussian DP [4]	GS	Light	(ε, δ)	×	×
Piecewise Laplace [3]	LS	Light	$(\varepsilon, 0)$	×	✓
ON–ALM (ours)	LS (tang.)	Light	(ε, δ)	✓	✓

vs. Smooth Sensitivity. The smooth sensitivity framework [2] achieves instance-adaptivity but requires heavy-tailed noise (Student-t or Cauchy) for pure ε-DP, resulting in infinite variance. ON–ALM uses light-tailed Laplace noise under (ε, δ)-DP with bounded variance, at the cost of a small additive δ and the feasibility constraint on β.

vs. K-Norm Gradient. The K-norm mechanism [5] aligns noise with a convex body derived from the query geometry but calibrates to a fixed global sensitivity and does not exploit spatial variation. ON–ALM additionally tracks how sensitivity varies across the manifold via the log-Lipschitz condition, providing strictly finer calibration for queries with non-uniform sensitivity.

vs. Gaussian DP. Gaussian DP [4] uses a global ℓ_2-sensitivity calibration and does not exploit the Riemannian structure. For queries on curved manifolds (e.g., covariance matrices on $\mathcal{S}^d_{++}$), the geodesic distance induced by the affine-invariant metric can be significantly smaller than the Euclidean distance, so Gaussian DP overestimates the required noise. ON–ALM uses intrinsic geodesic sensitivity, yielding a tighter noise calibration in these settings.

5.2. When ON–ALM is less suitable

ON–ALM is less suitable than competing mechanisms in three regimes: (i) *high output dimension* ($k \gg 10$), where the feasibility constraint becomes restrictive (see Section 5 and Table 1); (ii) *pure ε-DP requirements*, where the (ε, δ) relaxation is unacceptable; and (iii) *queries with unbounded Hessian*, where the log-Lipschitz condition cannot be verified and the stability parameter β may be difficult to estimate. In these cases, the smooth sensitivity framework or the piecewise Laplace mechanism [3] may be more appropriate.

6. Analytical illustration

To ground the theoretical framework in concrete computations, we provide two closed-form illustrations of ON–ALM that directly connect the mechanism parameters to the privacy-utility trade-off. These illustrations are exact and fully reproducible from the mechanism formula; empirical evaluation on real datasets is identified as a priority direction for future work (Section 7).

6.1. Illustration 1: Quadratic query on the euclidean disk

Let $M = \{x \in \mathbb{R}^2 : \|x\|_2 \leq 1\}$, $f(x) = \frac{1}{2}x^\top Ax$, $A = \mathrm{diag}(3, 0.5)$. The local sensitivity is $\Delta^* f(x) = \|Ax\|_2$ and the global sensitivity is $\mathrm{GS} = \|A\|_{\mathrm{op}} = 3$. With parameters $k = 1, \beta = 0.03, \xi = 0.01, \delta = 10^{-5}$, the mechanism overhead factor is

$$\mathrm{overhead}(\varepsilon) = \frac{e^{2\beta}\,\varepsilon}{\varepsilon - k\beta - \Lambda_{k,\delta}(e^\beta - 1)},$$

which equals 1.714 at $\varepsilon = 1$. The mechanism outperforms Lap-GS pointwise whenever $\Delta^*(x) < \mathrm{GS}/\mathrm{overhead}(\varepsilon)$; at $\varepsilon = 1$ this threshold is 1.750, holding over 69.3% of the unit disk (computed analytically). Table 3 reports exact noise scales and mean squared error (MSE) $= 2b^2$ at three representative points.

Table 3. Illustration 1: exact noise scale b and MSE $= 2b^2$ at three points of the unit disk. $x_\ell = (0.1, 0)$, $x_m = (0.5, 0)$, $x_h = (0.9, 0)$; Δ^* values 0.300, 1.500, 2.700. Parameters: $k = 1$, $\beta = 0.03$, $\xi = 0.01$, $\delta = 10^{-5}$. All values computed analytically from the mechanism formula.

	$\varepsilon = 1.0$		$\varepsilon = 2.0$		$\varepsilon = 3.0$	
Query point	b	MSE	b	MSE	b	MSE
Lap-GS (all x)	3.000	18.00	1.500	4.500	1.000	2.000
ON-ALM at x_ℓ	0.514	0.529	0.197	0.078	0.122	0.030
ON-ALM at x_m	2.572	13.23	0.984	1.936	0.608	0.740
ON-ALM at x_h	4.629	42.85	1.770	6.269	1.095	2.396
ON-ALM mean (disk)	2.265	13.60	0.866	1.990	0.536	0.761
Win fraction	69.3%		86.4%		91.1%	

The table shows that ON–ALM provides large pointwise gains in low-sensitivity regions (x_ℓ: 97% MSE reduction at $\varepsilon = 1$) while behaving as expected near the worst-case boundary (x_h: noise exceeds Lap-GS because $\Delta^*(x_h) \approx \mathrm{GS}$). The win fraction grows monotonically with ε, consistent with the overhead analysis.

6.2. Illustration 2: Trace query on the SPD manifold

Let $M = \mathcal{S}^3_{++}$ with the affine-invariant metric $g_P(U, V) = \mathrm{tr}(P^{-1}UP^{-1}V)$, $f(P) = \mathrm{tr}(WP)$, $W = \mathrm{diag}(1, 0.5, 0.1)$. For unit-trace diagonal $P = \mathrm{diag}(d_1, d_2, d_3)$ (normalized sample covariance constraint

$\text{tr}(P) = 1$), the local sensitivity reduces analytically to

$$\Delta^* f(P) = \|P^{1/2} W P^{1/2}\|_F = \sqrt{\sum_{i=1}^{3} (d_i w_i)^2},$$

and the global sensitivity is $\text{GS} = \|W\|_{\text{op}} = 1.0$. Table 4 reports exact values at three diagonal unit-trace test points with parameters $k = 1, \beta = 0.02, \xi = 5 \times 10^{-4}, \delta = 10^{-5}$.

Table 4. Illustration 2: trace query on $\mathcal{S}^3_{++}$. Points: $P_\ell = \text{diag}(0.1, 0.2, 0.7)$ ($\Delta^* = 0.158$), $P_m = \text{diag}(0.5, 0.3, 0.2)$ ($\Delta^* = 0.522$), $P_h = \text{diag}(0.9, 0.08, 0.02)$ ($\Delta^* = 0.901$). GS = 1.0. All values computed analytically.

	$\varepsilon = 1.0$		$\varepsilon = 2.0$		$\varepsilon = 3.0$	
Query point	b	MSE	b	MSE	b	MSE
Lap-GS (GS = 1.0)	1.000	2.000	0.500	0.500	0.333	0.222
ON-ALM at P_ℓ	0.220	0.097	0.094	0.018	0.060	0.007
ON-ALM at P_m	0.727	1.058	0.311	0.194	0.198	0.078
ON-ALM at P_h	1.255	3.148	0.537	0.576	0.341	0.233

At P_ℓ (weight concentrated on the low-w dimension), ON–ALM achieves 95.2% MSE reduction relative to Lap-GS at $\varepsilon = 1$. At P_h (near the worst-case configuration), ON–ALM is 57% worse than Lap-GS, consistent with the overhead exceeding the local-to-global sensitivity ratio. Both illustrations are exact consequences of the mechanism formula and serve to validate the theoretical predictions of Section 5.

7. Conclusions

We introduced a differential-geometric framework for privacy-preserving data analysis based on the directional local sensitivity operator Δ^*, defined on the tangent bundle of a Riemannian data manifold. Unlike classical global sensitivity, Δ^* isolates maximal directional derivatives, providing an instance-adaptive noise calibration that reflects local variability.

Using this operator, we developed the ON–ALM and proved its (ε, δ)-differential privacy guarantees. Because the regularized sensitivity satisfies a log-Lipschitz stability condition, finite-variance Laplace noise is sufficient under a typical-set analysis. This avoids the heavy-tailed noise required in pure ε-DP and couples the noise geometry directly to the ℓ_1 operator norm of Δ^*.

The structured comparison in Section 5 and the analytical illustrations in Section 6 demonstrate that ON–ALM achieves strictly lower noise scales than the global Laplace mechanism in regions of low local sensitivity, and maintains finite variance in contrast to smooth sensitivity constructions with Cauchy noise. The feasibility analysis in Section 5 (Table 1) delineates the regime of applicability: ON–ALM is best suited for scalar or low-dimensional outputs ($k \leq 5$–10) under moderate privacy budgets ($\varepsilon \geq 1$), and the three remedies discussed therein (dimensionality reduction, block composition, alternative noise geometry) provide a roadmap for extending its applicability to higher-dimensional settings.

Overall, this framework demonstrates how local geometric regularity yields stable, instance-adaptive privacy mechanisms. By establishing a mathematically rigorous foundation for noise calibration on manifolds, ON–ALM bridges the gap between theoretical differential privacy and the practical demands of privacy-preserving computational intelligence. For next-generation data science applications, where complex representations often naturally reside on non-Euclidean spaces, aligning the privacy mechanism with the underlying data geometry is essential for maintaining analytical utility without sacrificing formal guarantees.

Future work will pursue four directions. *Theoretical extensions*: (i) sharper dimension-dependent bounds via Rényi-DP composition; (ii) hybrid Riemannian–Gaussian mechanisms that inherit the light-tailed properties of ON–ALM while admitting higher output dimensions; (iii) fully data-driven estimation of β and ξ via differentially private Hessian approximation. *Empirical validation*: (iv) systematic experimental evaluation on real manifold-valued datasets, including SPD covariance streams in neuroimaging (diffusion tensor imaging, electroencephalogram-based brain-computer interface (EEG-BCI)), graph neural network embeddings on citation networks, and hyperbolic word embeddings in natural language processing (NLP). For such tasks, the present framework provides the theoretical foundations; the empirical evaluation requires careful protocol design (adjacency definition, sensitivity estimation, baseline implementation) that we identify as a dedicated research program. The analytical illustrations in Section 6 provide the quantitative predictions that such experiments should aim to validate or falsify.

Future work will include a detailed experimental comparison between the proposed ON–ALM mechanism and recent manifold-aware differential privacy approaches, including differentially private Riemannian optimization [6], private federated manifold learning [8], and conformal differential privacy mechanisms [10].

Use of AI tools declaration

The author declares he has not used Artificial Intelligence (AI) tools in the creation of this article.

Conflict of interest

The author declares there is no conflict of interest.

References

1. C. Dwork, F. McSherry, K. Nissim, A. Smith, Calibrating noise to sensitivity in private data analysis, in *Theory of Cryptography Conference (TCC)*, Springer, Berlin, Heidelberg, (2006), 265–284.
2. K. Nissim, S. Raskhodnikova, A. Smith, Smooth sensitivity and sampling in private data analysis, in *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, (2007), 75–84.
3. D. Durfee, Unifying Laplace mechanism with instance optimality in differential privacy, preprint, arXiv:2505.02798.
4. J. Dong, A. Roth, W. J. Su, Gaussian differential privacy, *J. R. Stat. Soc. B*, **84** (2022), 3–37. https://doi.org/10.1111/rssb.12454
5. M. Reimherr, J. Awan, KNG: The K-norm gradient mechanism, in *Advances in Neural Information Processing Systems*, **32** (2019).

6. A. Han, B. Mishra, P. Jawanpuria, J. Gao, Differentially private Riemannian optimization, *Mach. Learn.*, **113** (2024), 1133–1161. https://doi.org/10.1007/s10994-023-06508-5
7. C. Li, S. Long, H. Liu, Y. Choi, H. Sekiya, Z. Li, Enhancing sparse mobile crowdsensing with manifold optimization and differential privacy, *IEEE Trans. Inf. Forensics Secur.*, **19** (2024), 6070–6083. https://doi.org/10.1109/TIFS.2024.3407668
8. Z. Huang, W. Huang, P. Jawanpuria, B. Mishra, Federated learning on Riemannian manifolds with differential privacy, preprint, arXiv:2404.10029.
9. S. Garg, V. Torra, Privacy in manifolds: Combining k-anonymity with differential privacy on Fréchet means, *Comput. Secur.*, **144** (2024), 103983. https://doi.org/10.1016/j.cose.2024.103983
10. P. He, L. Tang, M. A. Rahimian, J. Joshi, Conformal-DP: Differential privacy on Riemannian manifolds via conformal transformation, preprint, arXiv:2504.20941.
11. M. P. do Carmo, *Riemannian Geometry*, Birkhäuser, Boston, 1992.
12. J. M. Lee, *Introduction to Riemannian Manifolds*, 2nd edition, Springer, Cham, 2018.
13. G. Casella, R. L. Berger, *Statistical Inference*, 2nd edition, Thomson Learning, Belmont, CA, 2002.
14. S. Dasgupta, A. Gupta, An elementary proof of a theorem of Johnson and Lindenstrauss, *Random Struct. Algorithms*, **22** (2003), 60–65. https://doi.org/10.1002/rsa.10073
15. C. Dwork, G. N. Rothblum, S. Vadhan, Boosting and differential privacy, in *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, (2010), 51–60.