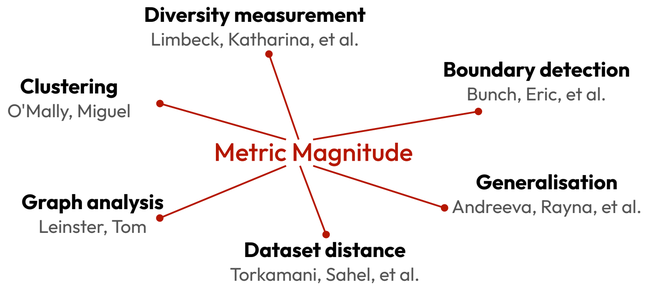
Privacy Requires Slicing

Differentially Private Magnitude via Projections

Sahel Torkamani · Henry Gouk · Rik Sarkar
School of Informatics, University of Edinburgh

Homepage

MOTIVATION



Magnitude has unbounded sensitivity, no fixed noise scale can guarantee differential privacy.

BACKGROUND

Magnitude of Metric Space

For a finite metric space, (X, d), we define the **similarity matrix** as $\zeta_X(x_i, x_j) := \exp(-d(x_i, x_j))$, for every $x_i, x_j \in X$.

A weighting of (X, d) is a function $w_X : X \to \mathbb{R}$ satisfying

$$\sum_{j \in X} \zeta_X(x_i, x_j)\, w_X(x_j) = 1$$

for every $x_i \in X$, where $w_X(x_i)$ is called the **magnitude weight**.

The **magnitude** of (X, d) is defined as $Mag(X, d) = \sum_{x_i \in X} w_X(x_i)$.

When X is a finite subset of $\mathbb{R}^D$, then ζ_X is invertible and magnitude is the sum of all the entries of the inverse of the similarity matrix.

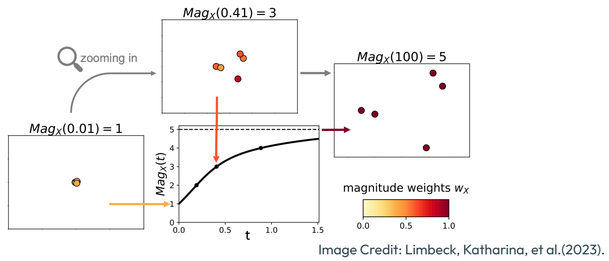


Image Credit: Limbeck, Katharina, et al.(2023).

Magnitude Function

For scaling parameter $t \in \mathbb{R}_+$, the **scaled metric space** (tX, d_t) is the metric space with the same points as X and metric $d_t(x, y) = t \cdot d(x, y)$.

The **magnitude function** assigns each finite metric space X to a family of scaled metric spaces $\{tX\}_{t>0}$ by $Mag_X(t) = Mag(tX)$.

THIS WORK

Sliced Magnitude

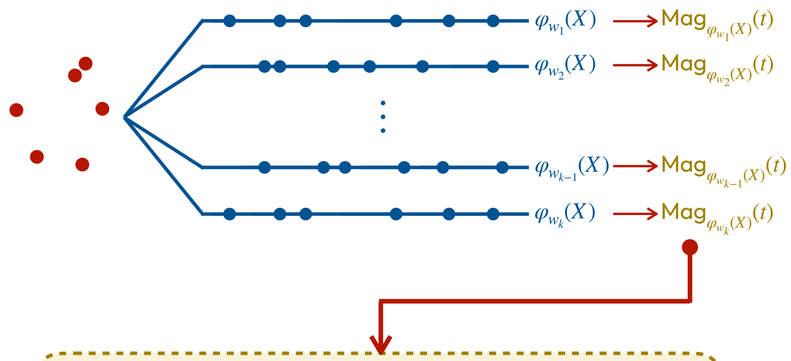
- **Idea:** aggregate magnitude over random 1-D projections of X, where it has a closed form — inspired by sliced Wasserstein.

Definition

Let $(\Omega, \mathcal{F}, \Phi)$ be a probability space and let $\varphi : \mathbb{R}^D \times \Omega \to \mathbb{R}$ be a measurable function. For each $w \in \Omega$, denote $\varphi_w(x) = \varphi(x, w)$. For a finite set $X \subset \mathbb{R}^D$, the sliced magnitude with scale parameter $t \in \mathbb{R}_+$ is defined as

$$\mathrm{SMag}_X(t) = \mathbb{E}_{w \sim \Phi}\left[\mathrm{Mag}_{\varphi_w(X)}(t)\right],$$

with sliced weighting as $\tilde{\mathbf{w}}^t_X = \mathbb{E}_{w \sim \Phi}\left[\mathbf{w}^t_{\varphi_w(X)}\right]$.

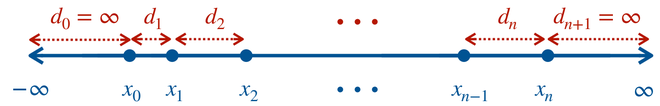


Corollary 2.3.4 (Leinster 2013) Every finite subspace of $\mathbb{R}$ has Mobius inversion. If $X = \{x_0 < \cdots < x_n\} \subseteq \mathbb{R}$ then, writing $d_i = x_i - x_{i-1}$,

$$\mathrm{Mag}_X(t) = 1 + \sum_{i=1}^{n} \tanh\frac{td_i}{2}.$$

- It has a closed form magnitude and weight:

$$\mathbf{w}^t_X(x_i) = \frac{1}{2}\left(\tanh\frac{td_i}{2} + \tanh\frac{td_{i+1}}{2}\right).$$

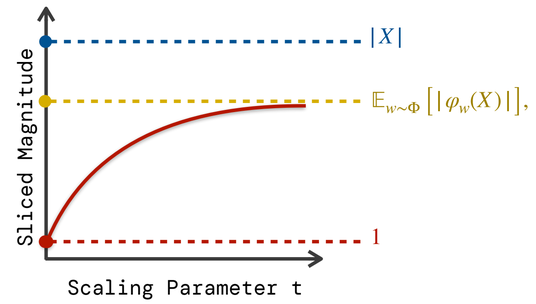


Sliced magnitude aggregates the magnitudes of 1-D projections of X, inheriting ℝ's good behavior while extending to any dimension D.

PROPERTIES

Behaviour Over Values of t

- **Small scale:** grows linearly, reflects expected projected diameter; approaches 1 as scale → 0.
- **Large scale:** converges to expected number of distinct projected points, bounded by set size.
- **Monotonic:** never decreases as scale increases.



DIFFERENTIAL PRIVACY

Differentially Private Sliced Magnitude

- **Idea:** We proved the sensitivity of $\mathrm{SMag}_X(t)$ is strictly bounded independent of dimension D and scale t:

$$0 \le \Delta\mathrm{SMag}_X(t) < 1$$

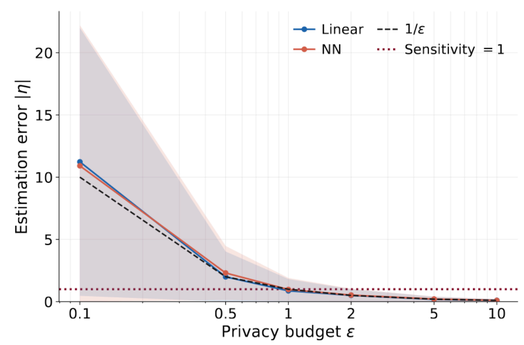
Theorem

- Let $\varepsilon > 0$ be the privacy parameter, releasing

$$\widetilde{\mathrm{SMag}}(X; t) = \mathrm{SMag}_X(t) + \mathrm{Lap}(1/\varepsilon)$$

is ε-differentially private.

- For (ε, δ)-DP with privacy parameter $\delta > 0$, the Gaussian mechanism with noise scale $\sigma = \sqrt{2\ln(1.25/\delta)}/\varepsilon$ applies.



sahel.torkamani@ed.ac.uk