Reconstruction and Secrecy under Approximate Distance Queries

NeurIPS 2025

Shay Moran — Technion, Google Research   Elizaveta Nesterova — Technion

Distance-query game — rules

1. **Responder** secretly chooses a target point $\mathrm{ref} \in X$.
2. For each round $i = 1, \ldots, T$:

   Reconstructor selects a query point q_i,

   Responder returns a noisy distance r_i.

Goal of the reconstructor

Output a final guess $\hat{x}_T$ that is as close as possible to ref.

Motivation

Perspective of the reconstructor.
GPS, ground-based navigation.

Perspective of the responder.
Protect sensitive data while still answering queries in an informative way.

Trade-off: *Reconstructor* seeks accuracy vs. *Responder* injects noise.

Problem setup

Noise parameters. Fix multiplicative ϵ and additive δ.

Game board. Work in a metric space (X, dist) and allow T rounds.

Roles.

- Reconstructor: chooses queries, wants small error.
- Responder: holds a hidden point ref and may add noise.

Problem setup

Noise parameters. Fix multiplicative ϵ and additive δ.

Game board. Work in a metric space (X, dist) and allow T rounds.

Roles.

Reconstructor: chooses queries, wants small error.

Responder: holds a hidden point ref and may add noise.

Each reply must be (ϵ, δ)–close to the true distance:

$$r \;\leq\; (1+\epsilon)\, d + \delta, \qquad d \;\leq\; (1+\epsilon)\, r + \delta.$$

After the T-th round

The reconstructor outputs its guess $\hat{x}_T$.

The responder reveals the hidden point ref.

Reconstruction error

$$\mathrm{err}_T \;=\; \mathrm{dist}_X\big(\hat{x}_T,\ \mathrm{ref}\big).$$

After the T-th round

The reconstructor outputs its guess $\hat{x}_T$.

The responder reveals the hidden point ref.

Reconstruction error

$$\mathrm{err}_T \;=\; \mathrm{dist}_X\big(\hat{x}_T,\ \mathrm{ref}\big).$$

Benchmark $\mathrm{OPT}(T, \epsilon, \delta)$

$$\mathrm{OPT}(T, \epsilon, \delta) \;=\; \inf_{\text{reconstructors}} \ \sup_{\text{responders}} \ \mathrm{dist}_X\big(\hat{x}_T,\ \mathrm{ref}\big).$$

OPT is the smallest worst-case error the optimal reconstructor can guarantee.

The quantity we track

$$\mathrm{OPT}(T, \epsilon, \delta), \qquad T = 1, 2, 3, \ldots$$

This minimax error is *non-increasing* in T.

We want to understand how fast it decreases.

The quantity we track

$$\mathrm{OPT}(T, \epsilon, \delta), \qquad T = 1, 2, 3, \ldots$$

This minimax error is *non-increasing* in T.

We want to understand how fast it decreases.

Long-run "optimal" value

$$\mathrm{OPT}_{\epsilon,\delta} := \lim_{T \to \infty} \mathrm{OPT}(T, \epsilon, \delta).$$

Question: how to describe the quantity $\mathrm{OPT}_{\epsilon,\delta}$ using the geometric properties of X?

Diameter–radius profile of a metric space X

Definition

For $\alpha \geq 0$, define

$$\mathsf{e}_X(\alpha) \;=\; \sup_{\substack{S \subseteq X \\ \operatorname{diam}(S) \leq \alpha}} \operatorname{radius}\big(\mathrm{MEB}(S)\big),$$

where $\mathrm{MEB}(S)$ denotes the *minimum enclosing ball* of S.

Intuition. How “spread out” can a set of diameter α be?

Euclidean space $\mathbb{R}^n$: $\mathsf{e}_X(\alpha) = \sqrt{\frac{n}{2(n+1)}}\,\alpha$.

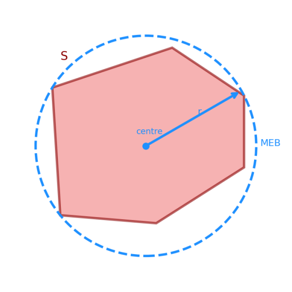

Red set S and its dashed MEB

Main Result I — Geometric Characterization

Theorem (Tight error via Chebyshev radius)

For every totally bounded metric space X,

$$\mathrm{OPT}_X(\epsilon, \delta) = \mathsf{e}_X\big((2+\epsilon)\delta\big).$$

Moreover, whenever distance α is realizable in X,

$$\alpha/2 \le \mathsf{e}_X(\alpha) \le \alpha.$$

Main Result II — Pseudo-Finiteness

Definition

A space X is *(ϵ, δ)-pseudo-finite* if

$$\exists T_0 : \ \mathrm{OPT}_X(T, \epsilon, \delta) = \mathrm{OPT}_X(\epsilon, \delta) \quad \forall\, T \geq T_0.$$

Main Result II — Pseudo-Finiteness

Definition

A space X is (ϵ, δ)*-pseudo-finite* if

$$\exists T_0 : \; \mathrm{OPT}_X(T, \epsilon, \delta) = \mathrm{OPT}_X(\epsilon, \delta) \quad \forall\, T \geq T_0.$$

Theorem (Pseudo-finiteness in convex Euclidean spaces)

Let $X \subset \mathbb{R}^n$ be bounded, convex, and $\dim X > 0$. Then for $\delta \ll \operatorname{diam} X$ the space X is not (ϵ, δ)-pseudo-finite, except for the trivial case $\epsilon = 0$, $\dim X = 1$.

Main Result II — Pseudo-Finiteness

Definition

A space X is (ϵ, δ)*-pseudo-finite* if

$$\exists T_0 : \ \mathrm{OPT}_X(T, \epsilon, \delta) = \mathrm{OPT}_X(\epsilon, \delta) \quad \forall T \geq T_0.$$

Theorem (Pseudo-finiteness in convex Euclidean spaces)

Let $X \subset \mathbb{R}^n$ be bounded, convex, and $\dim X > 0$. Then for $\delta \ll \mathrm{diam}\, X$ the space X is not (ϵ, δ)-pseudo-finite, except for the trivial case $\epsilon = 0$, $\dim X = 1$.

$\Rightarrow$ No finite number of queries can reach the optimum in higher-dimensional convex spaces.

Thank you!

Acknowledgements

Funded by the European Union (ERC, GENERALIZATION, 101039692). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Funded by the European Union