

Learning from Leakage: Database Reconstruction from Just a Few Multidimensional Range Queries

Peijie Li

Delft University of Technology

zimuqin1106@outlook.com

Huanhuan Chen[✉]

Delft University of Technology

h.chen-2@tudelft.nl

Evangelia Anna Markatou

Delft University of Technology

e.a.markatou@tudelft.nl

Kaitai Liang

University of Turku

kaitai.liang@utu.fi

Delft University of Technology

kaitai.liang@tudelft.nl

Abstract—Searchable Encryption (SE) has shown a lot of promise towards enabling secure and efficient queries over encrypted data. In order to achieve this efficiency, SE inevitably leaks some information, and a big open question is how dangerous this leakage is. While prior reconstruction attacks have demonstrated effectiveness in one-dimensional range query settings, extending them to high-dimensional datasets remains challenging. Existing methods either demand excessive query information (e.g., an attacker that has observed all possible responses) or produce low-quality reconstructions in sparse databases. In this work, we present REMIN, a new leakage-abuse attack against SE schemes in multi-dimensional settings, exploiting access and search pattern leakage from range queries. REMIN leverages unsupervised representation learning to transform query co-occurrence frequencies into geometric signals, enabling an attacker to infer relative spatial relationships among encrypted records. This approach allows accurate and scalable reconstruction of high-dimensional datasets under minimal leakage. Furthermore, we introduce REMIN-P, an active variant of the attack that incorporates a practical poisoning strategy. By injecting a small number of auxiliary anchor points, REMIN-P significantly improves reconstruction quality, particularly in sparse or boundary regions of the data space. We evaluate our attacks extensively on both synthetic and real-world datasets. Compared to state-of-the-art reconstruction attacks, our reconstruction attack achieves up to 50% reduction in mean squared error (MSE), all while maintaining fast and scalable runtime. Our poisoning attack can further reduce MSE by an additional 50% on average, depending on the poisoning strategy.

I. INTRODUCTION

Searchable encryption (SE) [2], [3] supports secure and efficient queries over encrypted databases. Current SE constructions incorporate various cryptographic techniques such as oblivious RAM (ORAM) [4], [5], fully homomorphic encryption (FHE) [6], and property-preserving encryption (PPE) [7] and they are widely implemented in real-world applications. For instance, MongoDB deploys Queryable Encryption (a

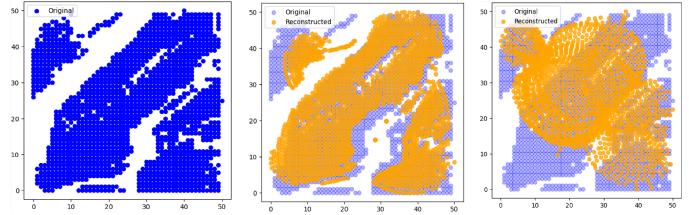


Fig. 1: (Left) Manhattan Highway Crossings dataset. (Middle) Reconstruction by our method REMIN, achieving an MSE of 14.25. (Right) Reconstruction by Markatou et al. [1], with an MSE of 42.97. (under 1% uniformly sampled queries)

variant of SE) [8] allowing users to perform expressive (e.g. range) queries on encrypted data.

In order to achieve its high efficiency, SE leaks some information to the server. Common types of leakage are *access pattern* (which records are returned given a query), *search pattern* (whether two queries are identical), and *volume pattern* (the number of records returned given a query). This leakage has been shown to allow a server to infer information about the plaintext data [9], [10], [11], leading to a line of research aiming to understand how this leakage can be exploited to reveal information about the encrypted data. Among various query types in SE, range queries pose a particularly severe leakage threat. Unlike keyword queries, which typically reveal only discrete match signals, range queries leak *structural information* by returning clusters of spatially or semantically adjacent records. When a query is issued, an associated token will be generated to retrieve records within a specified range as the response. As a result, attackers may recover layout characteristics of the original dataset, such as density clusters, relative positioning, or even record locations. In this work, we investigate the security risks of range queries across arbitrary dimensions under limited information leakage.

Motivation. A standard setting we consider here is a passive persistent model observes all issued queries and their encrypted responses. Under this threat model, prior work [12],

TABLE I: Comparison of our attack REMIN with closely related work. The “Required Leakage” column represents the number of queries needed for the attack, with fewer queries being more desirable. Both required leakage and reconstruction quality are indicated with stars, where more stars mean fewer queries required and better query efficiency. Our approach achieves practical and accurate database reconstruction using minimal leakage. Markatou et al. [1] achieves FDR in dense databases, and achieves high quality ADR in non-sparse datasets.

Attacks	Query Type	Assumption					Attack		
		Query Dist.	Database	Search Pat.	Access Pat.	ADR	FDR	Required Leakage	Reconstruction Quality
Kellaris et al. [12]	1D Range	Uniform	General	✗	✓	✗	✓	★☆☆	★★★
Lacharité et al. [13]	1D Range	Agnostic	Dense	✗	✓	✓	✓	★☆☆	★★★
Grubbs et al. [14]	1D Range	Uniform	General	✗	✓	✓	✓	★☆☆	★★★
Kornaropoulos et al. [15]	1D Range	Agnostic	General	✓	✓	✓	✓	★★★	★★★
Falzon et al. [16]	2D Range	Agnostic	General	✓	✓	✗	✓	★☆☆	★★★
Markatou et al. [17]	2D Range	Agnostic	General	✓	✓	✓	✓	★☆☆	★★★
Markatou et al. [1]	d D Range	Agnostic	(Dense)	✗	✓	✓	(✓)	★☆☆	★☆☆
REMIN	d D Range	Agnostic	General	✓	✓	✓	✗	★★★	★☆☆

[13], [14], [15], [18], [19] have demonstrated search pattern and access pattern can achieve full or approximate database reconstruction in one-dimensional (1D) datasets, where simpler relationships among records make such attacks more feasible. These attacks have achieved notable success in this regard, providing a solid foundation for understanding the potential of leakage-abuse attacks from range query leakage. However, as we move from one-dimensional to two-dimensional and even higher-dimensional databases, the complexity of the relationships between records increases significantly. To support attacks on these more complex scenarios, researchers have proposed various strategies to address the challenges posed by intricate data relationships and the exponential growth in equivalent databases (databases with the same leakage) [1], [16], [17], [20].

Across all existing methods, the core attack principle remains consistent: use the minimum amount of leaked information (e.g., partial query patterns) and maximize by the quality of the reconstructed data. As summarized in Table I, attacks targeting 1D datasets have achieved this, enabling accurate reconstruction with minimal leakage. However, higher-dimensional datasets naturally introduce greater complexity. While some attacks attain high reconstruction accuracy, they often rely on strong and impractical assumptions, such as access to all possible queries [16] or access to all possible responses [17]. These methods require substantial leakage, making them impractical for real-world settings.

The state-of-the-art graph-based approach by Markatou et al. [1] represents the first significant attempt to use minimal leakage in high-dimensional settings by observing very few queries and reconstructing from access pattern leakage alone. This attack has significantly weaker assumptions than earlier methods [16], [17]. But it primarily recovers local ordinal relationships using graph drawing techniques and fails to capture structural features like clustering and relative positioning, especially in generic datasets (i.e., non-dense or

sparse datasets). Consequently, its performance degrades in more complex scenarios, which raises a fundamental research question:

Under limited multi-dimensional range search leakage, what is the best reconstruction for general databases that we can achieve?

New Perspective. In this work, we propose a novel and practical database reconstruction attack that operates under minimal leakage assumptions, named REMIN. Unlike prior work such as Markatou et al. [1], which explores reconstruction under low query leakage but focuses primarily on **topological** recovery (i.e., identifying neighboring records) via an unweighted graph that captures connectivity but does not account for finer-grained distance signals, our goal is more ambitious: we aim to recover both the **topological** and **geometric** structure of the original database.

By topological structure, we refer to the connectivity and neighborhood relationships among data points—i.e., who is close to whom. By geometric structure, we mean the relative distances and spatial arrangement of data points in the underlying data space.

Our method builds on the observation that records retrieved together in response to user queries are often semantically or structurally related in the original dataset. By measuring how frequently pairs of records co-occur in leaked queries, we can estimate their relative proximity in the data space. Using these co-occurrence frequencies, we construct a high-dimensional distance matrix of size $n \times n$, where n is the number of records, and each record is represented as an n -dimensional vector (described in Section IV).

We then apply unsupervised representation learning and dimensionality reduction techniques (e.g., t-SNE [21]) to embed the records into a low-dimensional space, reconstructing both the topology and geometry of the data in its original target dimensional space, while preserving the pairwise distance information as faithfully as possible.

A key challenge in this setting is edge distortion: there are fewer queries covering boundary regions than central ones, resulting in inaccurate positioning and compression of sparse areas. To address this, we propose REMIN-P, a new active attack variant that employs a poisoning strategy. We consider a more powerful attacker, that is able to inject carefully crafted anchor points—records with known or leaked positions—into the dataset. These anchors act as geometric references, correcting misalignment and improving layout fidelity, particularly in sparse or underrepresented regions. To the best of our knowledge, this is the first poisoning-based reconstruction strategy tailored to range query attacks.

We evaluate our attacks on both synthetic and real-world datasets, simulating practical scenarios under different query distributions. A brief comparison is summarized in Table I, and the full experimental results are provided in Section V. Our main contributions are:

- **A practical reconstruction attack under minimal leakage:** We introduce an unsupervised, co-occurrence-based attack that recovers both topological and geometric properties for general datasets using only limited information leakage. REMIN remains effective even in extremely low-leakage conditions, even on sparse datasets (See Fig. 1 for an example).
- **A novel poisoning attack:** We propose REMIN-P, the first active attack variant, where the attacker poisons the dataset, injecting anchor points to improve global alignment and correct distortions in geometric reconstruction.
- **Comprehensive experimental evaluation:** We demonstrate the effectiveness of our attacks across diverse datasets and leakage scenarios, showing significant improvements in reconstruction accuracy and runtime compared to prior work.

Our work shows how minimal co-occurrence leakage can be exploited to recover both how records are related (topology) and how they are spatially arranged (geometry). This highlights the need for stronger leakage mitigation techniques in searchable encryption systems.

II. RELATED WORK

Searchable Encryption (SE) has evolved significantly from its early keyword-search origins to support increasingly complex functionality like dynamic updates and expressive queries [3], [22], [23], [24], [25], [26], [27], [28], [29], [30], [31]. Early constructions primarily supported simple keyword queries [2], [32], but were later extended to support more expressive query types, including conjunctive, boolean, and range queries [25], [33], [34], [35], [36]. In particular, Dynamic Searchable Symmetric Encryption (DSSE) [30] enables updates and deletions on encrypted databases, making SE practical for real-world applications.

In order to be practical, SE schemes inevitably leak side-channel information during query execution. Several works have systematically characterized and quantified the leakage in SE protocols [10], [11], [12], [37], [38], [39], which has been shown to be particularly exploitable for reconstructing private

data through range query attacks [10], [11], [12], [13], [40]. Most work on dataset reconstruction from range query leakage can be broadly categorized into three methodological classes: (1) symbolic or algebraic attacks that rely on geometric properties and deterministic constraints, (2) statistical estimation methods that infer data distributions from frequencies information, and (3) graph-theoretic approaches that exploit co-occurrence patterns to recover structural information. Each class of methods operates under different leakage assumptions and achieves varying levels of reconstruction accuracy and generality.

Early work primarily focused on determining how much information about the query inputs can be recovered from leakage [11], [37], [41]. Subsequently, several algebraic attacks have been proposed to exploit range query leakage for recovering plaintext geometry, which mainly explored the reconstruction of the encrypted dataset under a strong leakage assumption—full access pattern leakage and knowledge of the query distribution. KKNO [12] is the first to formalize this setting by introducing a generic volume-based attack model that combines access patterns with volume leakage. By analyzing the frequency of record occurrences across query responses, KKNO is able to infer exact record positions. Lacharité et al. [13] extended this approach to Approximate Database Reconstruction (ADR) with partial inaccuracies in the inferred structure in exchange for significantly improved efficiency. Falzon et al. [16] further demonstrated that leaked information in 2D databases can also be exploited to achieve exact reconstruction. These symbolic or order-based attacks offer strong recovery guarantees, but require full access pattern leakage and knowledge of the query or data distribution.

To avoid such strong assumptions, later methods turned to statistical estimation. Instead of directly solving for positions, these approaches estimate the underlying distribution or topological layout of the data using query response patterns. Kornaropoulos et al. [15] introduced a support-size estimator that, notably, was the first to leverage search pattern leakage—a previously overlooked signal—to reconstruct the database without relying on knowledge of the underlying query or data distribution. Building on this direction, Markatou et al. [17] proposed a non-parametric estimator that aggregates access frequencies to support partial reconstruction in two-dimensional settings. While these methods reduce the required leakage compared to earlier techniques, they still assume visibility over a large fraction of the query workload and can be sensitive to dataset sparsity and distributional skew.

A complementary line of work adopts a graph-theoretic perspective, shifting focus from precise reconstruction to structural recovery. Grubbs et al. [18] introduced a method that uses volume leakage to reconstruct the database structure in 1D settings. While the approach targets full reconstruction, it is primarily effective on dense datasets, resulting in an ordering of records rather than detailed spatial layout. Markatou et al. [1] further proposed a framework that constructs a co-occurrence graph from access patterns and achieves database reconstruction using graph-drawing techniques. These meth-

ods succeed under minimal leakage—often access pattern or volume pattern alone—and generalize well to higher dimensions. However, they model co-occurrence as binary and ignore richer signals such as relative distance or frequency, limiting their geometric fidelity.

In contrast to symbolic reasoning and graph-based approaches, machine learning offers a new paradigm for capturing geometric relationships from leakage. While ML techniques have been widely applied in privacy attacks—such as side-channel and website fingerprinting [42], [43]—they remain largely unexplored in the context of database reconstruction from searchable encryption leakage to the best of our knowledge. In particular, the effectiveness of unsupervised learning for this task remains an open question. Our work bridges this gap by leveraging co-occurrence frequency as a proxy for spatial proximity, embedding records via manifold learning, and refining the reconstruction through alignment. This enables robust and efficient recovery even under sparse leakage and agnostic query distributions.

III. PRELIMINARIES

Basic Concepts. In a searchable encryption (SE) scheme, a user first encrypts their dataset and uploads the encrypted data to an untrusted server. Later, to perform a range query, the user generates a query token that encodes the desired range and sends it to the server. Using this token, the server searches over the encrypted dataset and returns the identifiers of the documents whose associated values fall within the specified range.

Let $\mathcal{F} = \{r_1, r_2, \dots, r_n\}$ denote the dataset belonging to the user, where each r_i is a record in the dataset. We consider a d -dimensional dataset where each record is a unique point in a discrete space:

$$S_d = [N_1] \times [N_2] \times \dots \times [N_d],$$

where $[N_i] = \{0, 1, \dots, N_i - 1\}$ defines the index range along dimension i . Each record r has a domain value $x = (x_1, x_2, \dots, x_d) \in S_d$.

Let $\mathcal{E} = \{er_1, er_2, \dots, er_n\}$ be the encrypted dataset corresponding to \mathcal{F} . A range query q is defined by a set of bounds for each dimension:

$$q = [a_1, b_1] \times [a_2, b_2] \times \dots \times [a_d, b_d],$$

where $a_i, b_i \in [N_i]$ and $a_i \leq b_i$ for each i . The server responds with all encrypted records within the corresponding hyperrectangle:

$$\text{Response}(t_q) = \{er_{j_1}, \dots, er_{j_q} \mid r_j \in q, \forall j \in [j_1, j_q]\},$$

where t_q is the token corresponding to the query q . See Table VII in Appendix A-A for notation details.

Threat Model. We consider a setting similar to that studied in prior work on range query SE schemes [1], [12], [13], [14], [15], [17], [18], [19], where an honest-but-curious server stores both the encrypted dataset and the encrypted query tokens, and retrieves all query locally. As a consequence,

the adversarial server inevitably observes: (1) a subset of the encrypted records and query tokens, and (2) the relationship between them—specifically, whether a given encrypted record matches a given query token (i.e., whether it falls within the queried range).

The information revealed through these observations corresponds to well-studied *access pattern* and *search pattern* leakages. The access pattern reveals the relationship between each encrypted record and each query token, i.e., which (encrypted) records match which queries. The search pattern indicates whether two query tokens correspond to the same query, allowing the attacker to eliminate duplicate query tokens and thus avoid redundant computations in subsequent distance calculations. Considering that real-world adversaries are rarely purely passive, we additionally include a natural extension: beyond passively observing, the adversary may poison the data with a limited number of additional records. With this capability, the attacker can use the poisoned records to refine and stabilize global database reconstruction.

Attack Model. Reconstruction attacks attempt to recover the underlying structure of the encrypted database based on the observed leakage. We categorize these attacks into two types:

- **Full database reconstruction (FDR).** This attack aims to recover the **exact** positions of all records in the database, thereby reconstructing the entire dataset with complete accuracy.
- **Approximate Database Reconstruction (ADR).** This reconstruction seeks to recover the fundamental **topological** structure of the dataset, without necessarily recovering the exact positions of all records. The focus is on preserving the overall layout, such as clustering and relative positioning, rather than achieving a perfect reconstruction.

Practical SSE deployments inherently leak search and access patterns for scalability, enabling the server to observe query repetition, response volumes, and result overlaps—the well-studied L1/L2 leakage [11], [37]. In this work, we investigate approximate reconstruction under standard leakage assumptions adopted in state-of-the-art range-query attacks [12], [20]. Our threat model assumes a semi-honest server in a cloud-hosted searchable encryption service. The server passively observes the encrypted query responses after a sufficient delay, without tampering or injecting queries. Within this setting, we focus on limited leakage, where only a small fraction (e.g., 1%) of access and search patterns are available to the attacker, representing the minimal leakage baseline used in most prior attacks and defenses. Specifically:

- Partial search pattern leakage: The attacker knows a subset of queries token $t_q \mid q \in \mathcal{Q}$ that were issued to the database.
- Corresponding access pattern leakage: For each observed query, the attacker also knows the set of encrypted records returned in the response, as shown in Fig. 2 as Observed Response.

The goal of the attacker is to recover as much information as possible about the original dataset, minimizing the recon-

struction error and preserving the topological and geometric structure of the data.

Dimension Reduction. In our REMIN attack, we generate a representation matrix that captures the relative positions of records in a high-dimensional space. To convert this high-dimensional matrix into a reconstruction that reflects the original spatial distribution, we utilize dimensionality reduction techniques. Dimensionality reduction is a process used to reduce the number of dimensions in a dataset while retaining as much information as possible about the original structure. This technique is widely used in various fields, such as data visualization and pattern recognition, to simplify complex, high-dimensional data while preserving essential relationships between data points [44], [45], [46].

In the context of database reconstruction, the choice of dimensionality reduction method is critical. Classical dimensionality reduction methods like Principal Component Analysis (PCA) [47] and Multidimensional Scaling (MDS) [48] are designed to preserve global relationships, making them less suitable for the task at hand. Specifically, PCA assumes a linear structure, while MDS relies on the assumption that global distance information is reliable. However, the distance matrix constructed from the co-occurrence frequency matrix F is likely to exhibit a nonlinear, high-dimensional manifold structure. In this case, global geometric relationships are not the most informative, and dimensionality reduction techniques should focus on recovering local manifold structures instead.

Notably, t-SNE (t-distributed Stochastic Neighbor Embedding) [21] is particularly well-suited for this scenario because it prioritizes local relationships, making it effective at uncovering the intrinsic low-dimensional geometry in high-dimensional data. By minimizing the divergence between the original and projected distance matrices, t-SNE ensures that records that are close in the high-dimensional space remain close in the lower-dimensional space. In contrast, methods like PCA and MDS, which focus on global structure, may fail to capture important local relationships, leading to less accurate reconstructions. Thus, we use t-SNE to balance local and global structures effectively.

IV. BASIC ATTACK

We present a new attack called REMIN that leverages machine learning, frequency-based co-occurrence analysis and dimensionality reduction techniques. The core insight behind this attack is that the co-occurrence frequencies of records in query responses can serve as a high-dimensional representation of their spatial relationships. By applying dimensionality reduction algorithms, we can recover the underlying geometry of the dataset in a lower-dimensional (target) space (e.g., 2D or 3D). The full procedure is summarized in Algorithm 1 and an overview of the attack pipeline is in Fig. 2.

A. Extracting Frequency-Based Distances

For the leakage in range query SE schemes, we consider both access pattern and search pattern leakage, as described in Section III. In particular, the attacker leverages the leakage

Algorithm 1 REMIN(*responses*)

Input: *responses*: A key-value store mapping each observed query token to the set of encrypted records in response to that token

Output: \hat{P} : Reconstructed point locations

```

/* Step 1: Construct pairwise co-occurrence distance
matrix, see Algorithm 2 */
1:  $D \leftarrow \text{GETDISTANCEMATRIX}(\text{responses})$ 
/* Step 2: Apply t-SNE to obtain an initial embedding */
2:  $\hat{P}_0 \leftarrow \text{tSNE}(D)$ 
/* Step 3: Global alignment using Procrustes analysis */
3:  $\hat{P} \leftarrow \text{PROCRUSTES}(\hat{P}_0)$ 
4: return  $\hat{P}$ 
```

to infer the response set for each query token, denoted by $\{\text{Response}(t_q) \mid q \in \mathcal{Q}\}$, where each response comprises the encrypted records matching the query.

We begin by constructing a *frequency matrix* F that captures the co-occurrence relationships between encrypted records—specifically, whether two records appear together in the response to the same query. Since different queries may yield identical responses, we resolve such ambiguities using the search pattern leakage. The resulting matrix F serves as the basis for computing distances between records in subsequent steps.

To be precise, the co-occurrence frequency $F(i, j)$ is defined as the number of query tokens where encrypted records er_i and er_j are both returned in the corresponding responses:

$$F(i, j) = \sum_{t_q, q \in \mathcal{Q}} \mathbf{1}(er_i \in \text{Response}(t_q) \wedge er_j \in \text{Response}(t_q)).$$

An example is shown in Fig. 2, where $F(0, 2) = 2$, indicating that encrypted records a and c co-occur in two query responses. This matrix inherently reflects the spatial relationships among records, as records that are spatially closer are more likely to appear together in the same query response, resulting in higher co-occurrence frequencies. To convert this frequency-based information into a distance relationship, we apply appropriate distance metrics.

Common choices for this transformation include reciprocal distance (which we adopt for its efficiency and experimental accuracy) [49], Euclidean distance (less suitable in this case, as it assumes a real geometric relationship rather than frequency-based), and Gaussian kernel distance [50]. Reciprocal distance, defined as:

$$D(i, j) = \frac{1}{1 + F(i, j)} \quad (1)$$

is particularly effective, as it ensures that smaller co-occurrence frequencies (indicating weaker associations) result in larger distances, while higher frequencies (indicating stronger associations) correspond to smaller distances. Empirical comparisons across multiple distance metrics and dimensionality reduction methods are in the full version [51].

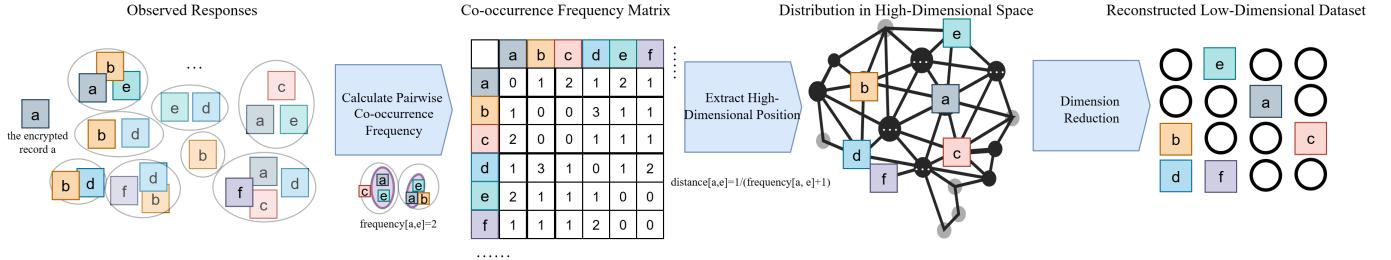


Fig. 2: A schematic illustration of the REMIN reconstruction attack. The method first observes the model’s responses to a set of query tokens and counts how frequently each pair of encrypted records (e.g., (a, b)) co-occurs in the different responses. Each cluster in the Observed Responses represents the response returned by a single query token, and each letter (e.g., a) corresponds to an individual encrypted record. These co-occurrence frequencies are transformed into pairwise distances (e.g., using the reciprocal of the frequency), resulting in a distance matrix that represents the records as points in a high-dimensional space. The high-dimensional relationships are then embedded into low-dimensional space through dimensionality reduction (e.g. t-SNE), preserving the structural relationships implied by record co-occurrence.

Algorithm 2 GETDISTANCEMATRIX(*responses*)

Input: *responses*: A key-value store mapping each observed query token to the set of encrypted records in response to that token

Output: *D*: Distance matrix where $D(i, j)$ is the dissimilarity measure

```

/* Get number of distinct encrypted records observed */
1:  $n \leftarrow |\bigcup_{(token, R) \in \text{responses}} R|$ 
/* Initialize the co-occurrence frequency matrix */
2: An  $n \times n$  symmetric matrix F  $\leftarrow \{\}$ 
3: for each  $(token, R) \in \text{responses}$  do
4:   for each pair of distinct records  $er, es \in R$  do
    /* Update frequency of point pair co-occurrence */
5:      $F[id_{er}, id_{es}] \leftarrow F[id_{er}, id_{es}] + 1$ 
6:      $F[id_{es}, id_{er}] \leftarrow F[id_{es}, id_{er}] + 1$ 
7:   end for
8: end for
/* Initialize distance matrix with infinity */
9: An  $n \times n$  symmetric matrix D initialized with all entries set to a large constant (e.g.,  $\infty$ )
10: for each  $(i, j) \in [n] \times [n]$  do
    /* Convert frequency to distance */
11:    $D(i, j) \leftarrow \frac{1}{1+F(i, j)}$ 
12: end for
13: for  $i \leftarrow 0$  to  $n - 1$  do
    /* Set distance to itself as zero */
14:    $D(i, i) \leftarrow 0$ 
15: end for
16: return D

```

The resulting distance matrix *D* as a proxy for the pairwise distances between records in the high-dimensional space, with its complete computation procedure detailed in Algorithm 2, where id_{er} denotes the index of an encrypted record *er*.

B. Dimensionality Reduction for Reconstruction

The distance matrix *D* captures the high-dimensional relationships between records based on their co-occurrence. However, the true underlying structure of the dataset exists in a lower-dimensional space. To recover this structure, we apply dimensionality reduction techniques.

Dimensionality reduction is a common machine learning approach used to simplify complex, high-dimensional data while preserving its essential features. In our method, we employ t-SNE, which is particularly effective at preserving local relationships and minimizing global distortions, ensuring that the reconstructed dataset reflects the original spatial layout. The result of this step is a set of reconstructed coordinates in a low-dimensional space, approximating the original dataset’s spatial relationships.

While t-SNE excels at preserving local structure, it is sensitive to parameters such as perplexity, which controls the trade-off between local and global preservation. For completeness, we refer to the full version [51] for a more detailed discussion on parameter sensitivity. Our empirical analysis shows that the optimal perplexity value scales with dataset size: smaller datasets benefit from lower perplexity values, while larger datasets require higher values for optimal reconstruction. By carefully selecting perplexity, we can significantly improve the reconstruction quality, balancing local and global structure preservation.

C. Refinement and Alignment

The dimensionality reduction process provides an initial set of coordinates for the records in the original space. However, due to the lack of absolute positional information in the co-occurrence data, the reconstruction result may suffer from potential misalignment. It typically differs from the original structure by certain rotations, scalings, and translations—affine transformations that preserve relative distances and angles between points but do not maintain their absolute positions.

To address this, we apply Procrustes analysis, a statistical method specifically designed to correct affine transformations.

Procrustes analysis requires a reference set of points for alignment, and we provide the detailed algorithmic description in the full version [51]. For datasets with a relatively uniform distribution, where records are evenly spaced across the spatial domain, we construct a reference grid to serve as the target for alignment. This approach works because each record in the original database lies on a grid point, making the grid a natural representation of the true spatial structure. Mathematically, given the original coordinates X (the reference grid) and the reconstructed coordinates Y , Procrustes analysis solves the following optimization problem:

$$\min_{R,s,t} \|X - (sYR + t)\|^2 \quad (2)$$

where R is a rotation matrix, s is a scaling factor, and t is a translation vector. This ensures that the reconstructed embedding aligns with the original dataset's spatial layout, correcting for reasonable distortions.

For datasets requiring integer coordinate constraints, simulated annealing, a probabilistic optimization technique, can be used to further optimize point placement, guaranteeing reconstructed coordinates occupy exact integer positions. Details on this technique are in our full version [51].

V. EVALUATION OF REMIN ATTACK

A. Experiment Setup

We assess the performance of our reconstruction method through comprehensive experiments across various datasets, including both synthetic and real-world scenarios. These datasets are as follows:

- **Grid (Synthetic):** A procedurally generated d -dimensional grid dataset (n^d) with uniform spacing and configurable sparsity (by randomly removing records). The grid size, denoted by n , indicates the number of cells per dimension. This dataset is used to simulate idealized structured data. *We refer to this dataset as Grid.*
- **California Intersections:** A real-world spatial dataset of over 21,000 road intersections in California, representing urban and suburban sprawl (normalized to a 50×50 grid), previously used in SE research [1], [17]. *We refer to this dataset as Cali.*
- **Amsterdam Drinking Water Points:** A spatial dataset of water access locations in Amsterdam, collected from OpenStreetMap. It exhibits moderate clustering typical of urban infrastructure. *We refer to this dataset as AMS.* The data was normalized to a 50×50 grid with 289 points.
- **Manhattan Highway Crossings:** This dataset captures highway intersections in Manhattan with 1708 points normalized to 50×50 space, exhibiting prominent spatial clustering. *We refer to this dataset as Manhattan.*
- **Paris Shops:** Locations of shops in central Paris, characterized by non-uniform spatial clusters that reflect commercial hotspots. *We refer to this dataset as Paris.* The data was normalized to a 50×50 grid with 1158 points.

- **Shanghai Bus Stops:** Bus stop locations in downtown Shanghai with 1046 points normalized to 50×50 , forming a relatively uniform distribution shaped by the urban road network, with non-equidistant spacing between points. *We refer to this dataset as Shanghai.*

- **New Hampshire Elevation (3D):** A 3D terrain dataset containing elevation samples from the White Mountains, previously used in reconstruction research [1]. Normalized to $16 \times 16 \times 16$. *We refer to this dataset as NH.*

To ensure a fair comparison between reconstruction methods, we align all reconstructed coordinates to the original dataset before computing evaluation metrics. This is necessary because reconstruction from range queries is inherently non-unique: many databases can produce identical query responses, forming what is known as the *reconstruction space* [1]. Within this space, any dataset is considered indistinguishable from others based solely on query leakage. To account for this ambiguity, we apply Procrustes analysis using known point correspondences to align reconstructions to the ground truth, allowing only rotation and scaling. This alignment ensures that evaluation metrics such as MSE and neighborhood accuracy reflect meaningful geometric fidelity rather than coordinate frame discrepancies. Unlike the basic attack setting, where alignment and simulated annealing are performed post-reconstruction, we avoid such post-processing as it may introduce artificial artifacts that could distort the comparison. All quantitative metrics reported in this section are computed after this alignment, ensuring that the results are directly comparable across methods and settings.

We evaluate our method using four key metrics, each capturing a distinct aspect of reconstruction quality:

- **Mean Squared Error (MSE):** MSE measures the average squared distance between the reconstructed coordinates and the original coordinates across all records. It quantifies the overall reconstruction error, balancing the contributions of both large and small errors, thus offering a comprehensive measure of reconstruction quality.
- **Tolerant Match Rate:** This relaxed version of exact coordinate match assesses the method's ability to recover positions within a specified radius of the ground-truth location. Specifically, a predicted point is considered a match if its Euclidean distance to the true point is less than $\sqrt{2}$. This metric is particularly valuable in practical settings, as it accounts for minor distortions that remain even after alignment, providing a more meaningful measure of localization accuracy.
- **Neighbor Accuracy:** This metric evaluates the preservation of local consistency by measuring the percentage of records whose 5 nearest neighbors in the reconstructed dataset match those in the original dataset. This is crucial for assessing how well the method preserves clusters and neighborhood relationships, which are often key to the dataset's structure.
- **Chamfer Distance:** Chamfer Distance measures the average closest-point distance between two sets of points,

capturing the overall structural similarity between them. Unlike MSE, which is sensitive to small shifts in coordinates, Chamfer Distance is robust to minor mismatches in point correspondence, making it ideal for evaluating the structural integrity of the reconstructed dataset.

These four metrics collectively evaluate both global accuracy and local consistency, providing comprehensive assessment of reconstruction quality. To account for the inherent instability of t-SNE and ensure statistical reliability, we repeat each experimental setup ten times in our experiments, and the results are averaged across these runs.

We note that our attack is domain-agnostic and depends only on the underlying data distribution. The geospatial datasets described above already span diverse structural patterns (dense, sparse, uniform). To further verify generality, we evaluate representative non-geospatial datasets (medical, financial, environmental [52], [53], [54]), demonstrating comparable reconstruction performance (see Appendix A-B).

B. Structural Visualization on Real-World and 3D Datasets

To qualitatively assess the reconstruction capabilities of our method, we present visual comparisons between the ground truth datasets, our reconstructed outputs, as well as those of a leading query-efficient baseline (Markatou et al. [1]). These comparisons span both real-world 2D spatial datasets obtained from OpenStreetMap and a 3D dataset, showing that our method outperforms the baseline in both geometric and topological preservation under severely limited query access.

We evaluate three real-world 2D datasets described in Section V-A: Amsterdam drinking water access points, Manhattan highway crossings, and shops in central Paris. All reconstructions are conducted with only 1% query coverage sampled under a uniform query distribution.

Fig. 1 and Fig. 3 depict datasets with strong clustering structures (AMS and Manhattan). Our method effectively recovers both the intra-cluster density and inter-cluster spacing, preserving neighborhood relationships with high fidelity. In the AMS dataset, we achieve 82% neighbor accuracy and a Chamfer distance of 1.46. Similar performance is observed for Manhattan. In contrast, the baseline [1] fails to capture both the global and local structures: on AMS it yields only 51% neighbor accuracy, and on Manhattan it reaches 55%. Its reliance on unweighted connectivity producing nearly uniform point clouds with significantly lower neighbor accuracy and substantially higher Chamfer distances, indicating both geometric and topological distortions.

Not all spatial datasets exhibit strong clustering. To evaluate reconstruction performance in general cases, we consider the Paris datasets (Fig. 4), which features more diffuse yet structurally meaningful point distributions. Despite the absence of large-scale clustering, this dataset still contains significant spatial heterogeneity and alignment patterns. Our method successfully recovers the relative positioning of points and underlying spatial organization. The reconstruction achieves consistent accuracy with 80.03% neighborhood accuracy and

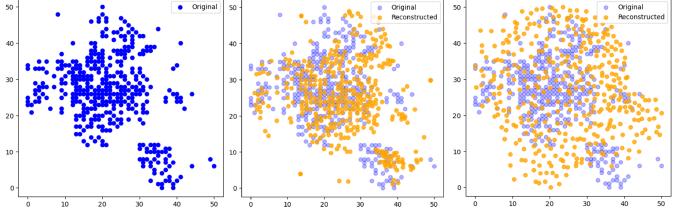


Fig. 3: (Left) Original Amsterdam Drinking Water dataset (50×50 grid). (Middle) Reconstruction by REMIN using 1% uniformly sampled queries, achieving 82.00% neighbor accuracy, Chamfer distance of 0.78, and MSE of 23.52. (Right) Reconstruction by the method of Markatou et al. [1] under the same setting, achieving 50.92% neighbor accuracy, Chamfer distance of 1.05, and MSE of 37.86.

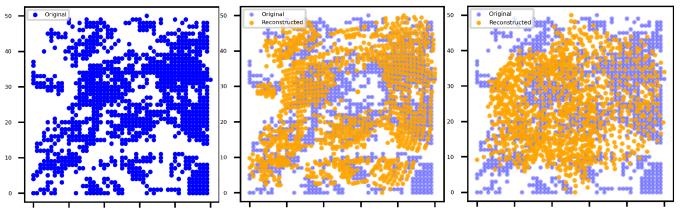


Fig. 4: (Left) Original Paris Shops dataset (50×50 grid). (Middle) Reconstruction by REMIN using 1% uniformly sampled queries, achieving 80.03% neighbor accuracy, Chamfer distance of 0.81, and MSE of 41.1. (Right) Reconstruction by the method of Markatou et al. [1] with 30% neighbor accuracy, Chamfer distance of 1.21, and MSE of 128.4.

low MSE, demonstrating robustness beyond clustered scenarios. Conversely, the baseline [1] once again fails to preserve spatial structure under the same query budget, producing a near-uniform random distribution of points. Similar results are observed under biased query distributions for all datasets (Appendix A-B).

To demonstrate the dimensional agnosticism of our approach, we further evaluate our method on a 3D dataset (NH), embedded in \mathbb{Z}^3 . Using the same query rate of 1% under uniform query distribution, we reconstruct the spatial layout via our co-occurrence-based dimensionality reduction pipeline. As shown in Fig. 14, the reconstructed shape preserves most of the original **topological** structure, highlighting that our method is dimension-agnostic. Since the approach solely relies on inter-point distances rather than absolute positions, it naturally extends to any intrinsic dimensionality, as long as local neighborhood information is well represented.

C. Comparison with Prior Work

To quantitatively evaluate the effectiveness of our reconstruction method, we compare its performance against representative baselines across a variety of settings. We first consider the earlier approach of Markatou et al. [17], which requires access to a full ordering of the encrypted dataset and therefore requires access to a large fraction of query responses. These assumptions are incompatible with the low-

leakage regime studied in this work—particularly, the setting where only 1% of all query responses are observed—making the method inapplicable under our evaluation framework. For empirical comparison, we therefore focus on the more recent query-efficient baseline of Markatou et al. [1], which is specifically designed for limited-leakage scenarios. Specifically, we assess reconstruction quality and scalability on both synthetic and real-world datasets.

First, we analyze performance on synthetic grid datasets, varying both dataset size and sparsity level to investigate robustness under different density and completeness conditions. Second, we perform extensive comparisons on multiple real-world 2D spatial datasets, focusing on the performance at very low (1%) and moderate (25%) query ratios under uniform query sampling. These ratios represent scenarios where the attacker obtains either very limited information or relatively more, but still insufficient, data for full reconstruction. By testing these methods on practical datasets, we aim to better demonstrate their effectiveness and applicability in real-world settings.

These experiments provide a comprehensive assessment of accuracy, robustness, and scalability, highlighting the significant advantage of our method over existing techniques under realistic leakage constraints.

1) Synthetic Dataset Comparison: To further validate the effectiveness and generality of our method, we benchmark it against Markatou et al. [1]. Although this technique also aims to recover database structure under limited query access, it tends to overlook the underlying pairwise distance in datasets.

We then design two evaluation scenarios to capture realistic variations in database conditions:

- **Scalability:** Datasets may vary in size or resolution, ranging from small tables to large-scale records.
- **Data Density:** Real-world datasets often exhibit non-uniform density, with missing records or irregular data collection leading to sparsity and gaps.

To assess performance under these conditions, we conduct two comparative studies:

- **Varying grid size.**
- **Varying sparsity.**

As mentioned before, once the mapping between reconstructed results and original coordinates is established, we compute all evaluation metrics based on the matched pairs, as shown in Fig. 5.

Scalability. We evaluate the scalability of our approach against the baseline [1] by testing both methods on grid-structured datasets, ranging from 20×20 to 40×40 in size. The experiments maintain a fixed query ratio of 1% and explore varying data distributions: uniform, beta, and Gaussian. This setup isolates the impact of database scale on reconstruction performance. As dataset size increases, a robust method should maintain reconstruction accuracy while remaining computationally efficient. To fully characterize this trade-off, we augment our evaluation with runtime measurements alongside the previously introduced metrics(See Fig. 15).

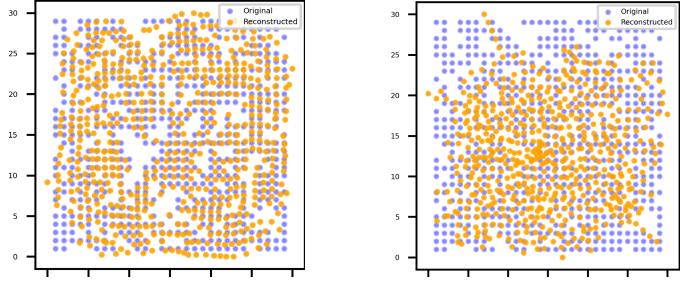


Fig. 5: Reconstruction by REMIN (Left) and by Markatou et al. [1] (Right) using 1% uniformly sampled queries.

The experimental results in Fig. 6 reveal that our method consistently outperforms the baseline [1] across all settings.

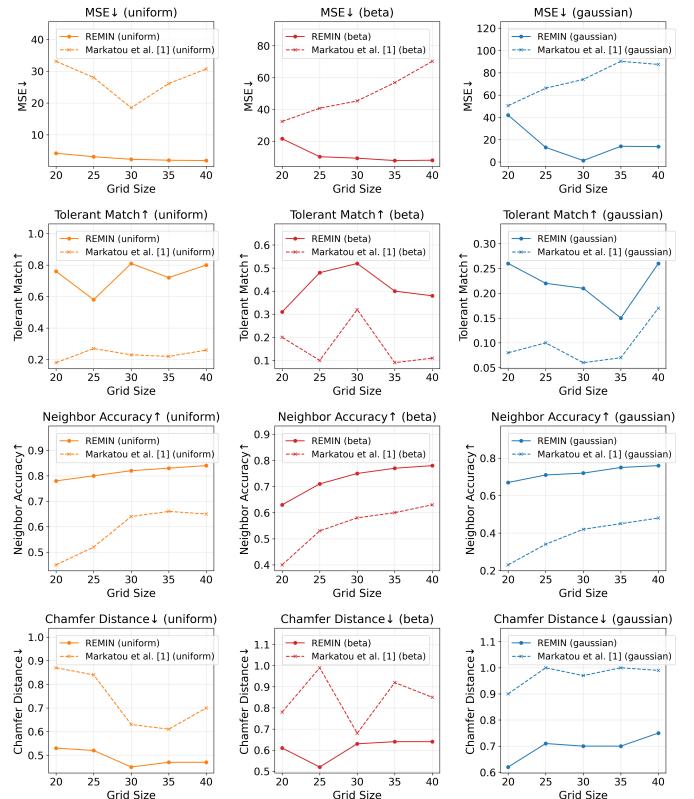


Fig. 6: Reconstruction performance on varying grid sizes (20×20 to 40×40) and data distributions with 1% query ratio.

Interestingly, as the grid size increases, both methods demonstrate stable performance, suggesting that both approaches exhibit scalability with respect to dataset size. However, we observe a more pronounced performance decline when varying the query distribution. Under non-uniform distributions, such as beta and Gaussian, both methods suffer in reconstruction quality due to increased sampling redundancy—fewer unique range responses result in less informative co-occurrence data. Nevertheless, our method maintains strong performance in preserving local structures. For example, neighborhood accuracy remains above 70% across all

settings, demonstrating that our method is particularly resilient to distributional shifts. While other metrics, such as tolerant match rate and MSE, show greater variability, neighborhood accuracy remains relatively stable, indicating that our approach excels at capturing fine-grained local geometry, even when global information is limited.

In addition to improved accuracy, our method is significantly more efficient. Runtime comparisons (See Fig. 15 in Appendix A-C) demonstrate an order-of-magnitude speedup, making our method particularly suitable for large-scale applications where traditional cubic-complexity methods become infeasible.

Overall, these results highlight the scalability, geometric accuracy, and topological robustness of our method. It consistently outperforms the state-of-the-art across a variety of settings, demonstrating its ability to effectively handle larger datasets and reliably preserve local structure even in the presence of sparse or biased query patterns.

Data Density. In many real-world applications, databases are not fully populated—due to missing records, data corruption, or partial collection—leading to what we refer to as gaps or missing data points. Existing methods, such as Markatou et al. [1] do not explicitly preserve local distance consistency and thus experience significant distortion when applied to sparser datasets.

To evaluate the robustness of our approach under varying levels of sparsity, we simulate random missing data points on a 30×30 grid with a uniform data distribution. The query ratio is fixed at 1%, and we gradually increase the percentage of missing records from 0% to 70%. All other experimental parameters remain consistent with the previous setup.

As shown in Table II, the accuracy of both methods decreases with increasing sparsity, as expected, since query results contain fewer co-occurring pairs. In sparser datasets, each query covers a smaller portion of the data, leading to reduced overlap between responses. Our method consistently outperforms the baseline by Markatou et al. [1] across all sparsity levels. In the range of 0%–60% missing data, our reconstruction remains notably stable, suggesting that the method is resilient to moderate sparsity. For example, even with 60% of the records missing, our method maintains a tolerant match rate above 49%, indicating successful structural recovery despite substantial data loss. In contrast, the baseline method exhibits a clear decline in all performance metrics as the missing data ratio increases.

Interestingly, across all levels of sparsity, our method maintains high neighborhood accuracy, same as the previous experiment (exceeding 80%), highlighting its robustness in preserving local topology. This resilience stems from our design, which leverages co-occurrence frequency to encode proximity, even when full structural information is not available.

However, when the missing data ratio exceeds 70%, performance becomes more volatile. At this point, co-occurrence signals become too sparse for reliable inference, and the gap between our method and the baseline [1] narrows. This limitation underscores a known challenge in high-sparsity scenarios:

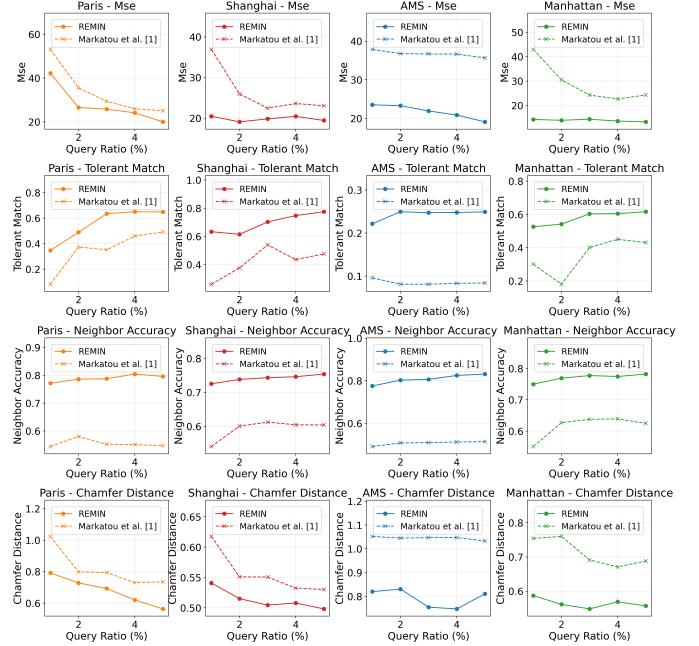


Fig. 7: Reconstruction performance on real-world datasets under varying query ratios (1% to 5%) with uniform query sampling. Our method consistently outperforms the method of Markatou et al. [1] in all metrics.

when neighbor relationships are insufficiently sampled, even distance-aware methods struggle to recover accurate structures.

In summary, the experiment demonstrates that our method is robust and scalable under realistic levels of sparsity (up to 60%), with clear advantages over the state-of-the-art. Its ability to preserve local geometry makes it well-suited for moderately sparse database scenarios, though extreme sparsity remains challenging for all current approaches.

2) Quantitative Comparison on Real-World 2D Spatial Datasets: In realistic scenarios, attackers often face strict constraints on the number of observable query responses, either due to system-imposed query budgets or privacy-preserving mechanisms. Consequently, a key challenge in reconstruction attacks is recovering as much spatial structure as possible with limited information. To assess the effectiveness of our method under such low-leakage conditions, we evaluate its performance on four real-world 2D datasets, and compare it against the state-of-the-art reconstruction baseline [1]. To simulate realistic query leakage, we uniformly sample a fixed percentage of range queries (from 1% to 5%) under a uniform distribution and use them as the sole input to each reconstruction method. We then evaluate the output using the four metrics mentioned before.

As shown in Fig. 7, the performance of all methods generally improves with higher query ratios, with MSE decreasing and accuracy-related metrics increasing. Across all datasets and query budgets, our method consistently outperforms the baseline [1] by a clear margin. Notably, under the lowest query ratio of 1%, our method achieves a tolerant match rate that

TABLE II: Comparison of reconstruction methods under increasing sparsity levels with perplexity set to 80.

Missing Data Ratio	REMIN				Markatou et al. [1]			
	MSE ↓	Tolerant Match ↑	Neighbor Accuracy ↑	Chamfer Distance ↓	MSE ↓	Tolerant Match ↑	Neighbor Accuracy ↑	Chamfer Distance ↓
0%	1.81	0.7365	0.8142	0.4375	30.10	0.3718	0.6898	0.5665
10%	2.19	0.7090	0.8103	0.4931	28.17	0.2761	0.6635	0.6035
20%	2.33	0.6715	0.8202	0.5115	28.25	0.1683	0.6096	0.6771
30%	2.44	0.6572	0.8335	0.5636	28.91	0.1146	0.5443	0.8047
40%	2.88	0.6251	0.8221	0.5951	23.59	0.1420	0.4951	0.7946
50%	10.33	0.4920	0.8132	0.6685	22.93	0.1178	0.4475	0.8575
60%	17.03	0.4109	0.8105	0.7983	22.65	0.1437	0.3944	1.0110
70%	18.79	0.2292	0.8071	0.9523	23.59	0.0859	0.3565	1.0540

more than doubles that of Markatou et al. [1]. Furthermore, due to the geometry-aware properties of our t-SNE-based reconstruction, our approach reliably preserves local spatial relationships, achieving neighbor accuracy above 72% across all settings even under minimal information observed.

As mentioned before, one of the most critical concerns is obtaining the highest quality of the reconstructed dataset under minimal query leakage. A practical attack should be capable of extracting meaningful structural information under extremely constrained leakage, while also scaling to higher coverage when available.

To further assess this, we fix the query ratio at two representative levels—1% and 25% (extremely low and relatively low query coverage)—and compare MSE results across the four real-world datasets.

The results, summarized in Table III and Table IV, demonstrate that our method is substantially more query-efficient than the baseline [1], achieving an average MSE reduction of approximately 50% across datasets. For instance, under just 1% query coverage, our method achieves up to 80% reduction in MSE relative to the baseline on Manhattan dataset, and in some cases achieves an MSE as low as 1.29.

These findings highlight the robustness of our framework in sparse regimes, demonstrating that even minimal query access can yield accurate geometric reconstructions—reinforcing the threat posed by range query leakage in real-world systems.

TABLE III: MSE of reconstruction using 1% query ratio.

Method	Paris	Shanghai	AMS	Manhattan
REMIN	42.17	22.81	23.52	14.25
Markatou et al. [1]	53.07	36.92	37.86	42.97

TABLE IV: MSE of reconstruction using 25% query ratio.

Method	Paris	Shanghai	AMS	Manhattan
REMIN	20.46	21.26	19.04	1.29
Markatou et al. [1]	26.14	23.02	37.20	10.82

D. Reconstruction in Higher Dimensions

As dimensionality grows, the reconstruction task becomes significantly more challenging due to sparser co-occurrence

observations and more complex geometric structures. We extend our evaluation to synthetic and real-world higher-dimensional datasets.

For the synthetic study, we compare our method against the leading query-efficient baseline [1] on grid-structured datasets with increasing dimensionality (2D to 6D) under a fixed query coverage of 1%, uniform query distribution, and 0% sparsity. To manage the exponential growth in data size and computational cost, we use progressively smaller grids (e.g., 16^3 for 3D, 6^5 for 5D). Evaluation is based on MSE and Chamfer distance, which captures geometric fidelity. We omit tolerant match and neighbor accuracy in dimensions ≥ 3 , as the geometric complexity and uniformity of inter-point distances in high-dimensional spaces render local neighborhood metrics less reliable.

TABLE V: Performance on higher-dimensional grid datasets with 1% query ratio under uniform distribution.

Dim	Method	MSE	Chamfer Dist
2D (32^2)	REMIN	0.9753	0.4113
	Markatou et al. [1]	24.6954	0.5184
3D (16^3)	REMIN	0.2991	0.5414
	Markatou et al. [1]	0.5968	0.5969
4D (8^4)	REMIN	0.6547	0.6851
	Markatou et al. [1]	0.6814	0.7061
5D (6^5)	REMIN	0.8624	0.9174
	Markatou et al. [1]	1.0639	1.0369
6D (4^6)	REMIN	1.1759	1.0247
	Markatou et al. [1]	1.1888	1.1168

As shown in Table V, our method consistently achieves lower MSE and Chamfer distance across all dimensions. Even in 6D, our attack maintains a clear advantage. This suggests that our distance-based co-occurrence representation preserves meaningful spatial structure even in higher dimensions.

We further vary sparsity levels (0–50%) across dimensions from 2D to 4D while keeping the total number of records fixed (64^2 , 16^3 , and 8^4). This isolates the effect of dimensionality and data density on reconstruction quality. Results in Table IX (Appendix A-C) indicate that reconstruction accuracy declines with increasing sparsity, consistent with the 2D observations.

Despite this trend, our method consistently yields lower MSE than the baseline [1] across all configurations. The performance gap remains stable across dimensions, showing that the co-occurrence-based distance formulation retains robustness under sparse query coverage even in higher-dimensional spaces.

We further validate the generality of our approach using non-geospatial real-world multidimensional datasets introduced earlier [52], [53], [54] by extending with one additional feature dimension selected from the original data attributes. Details of these datasets are provided in Appendix A-B. Evaluation under the same query coverage (1%) confirms consistent reconstruction behavior: our method consistently outperforms the baseline [1] across all three real-world datasets, yielding MSE reductions of approximately 20–60%. Specifically, reconstruction error decreases from 1.00 to 0.79 on Obesity, from 1.48 to 1.29 on Telecom, and from 2.17 to 0.92 on Landsat.

These findings demonstrate the scalability and robustness of our approach. Despite extreme query sparsity and dimensional complexity, our attack can still recover the structural layout of the dataset far more effectively than existing methods. This underscores the potential risks of leakage in high-dimensional encrypted systems, where traditional defenses may underestimate the attacker’s reconstruction capabilities.

E. Evaluation under Countermeasures

While SSE inevitably leaks side-channel information such as access and search patterns, real-world deployments often incorporate lightweight countermeasures to mitigate such leakage. Industrial systems and prototype implementations alike—such as ShieldDB [55], Veil [56] and CryptDB [57]—commonly employ techniques like query padding, randomized response sets, or canonical range partitioning to reduce the observability of sensitive correlations. These strategies trade additional bandwidth or computation for improved confidentiality, but their actual effectiveness against modern reconstruction attacks remains underexplored. To assess the real-world viability of REMIN under such realistic deployment conditions, we evaluate its robustness against representative leakage-mitigation defenses widely adopted in practical SSE systems. Specifically, we implement and test these countermeasures on a controlled 30×30 grid dataset with 1% uniformly sampled range queries.

We focus on two representative defenses. First, *max padding* [37], [58], [59], where responses are padded with fake records to the next power-of-two or a fixed multiple of x , mitigates leakage while offering a tunable privacy–cost trade-off. Second, *obfuscation with BT-style canonical coverage* [60], which returns results for the smallest precomputed canonical range whose span is a power of two, thereby obfuscating access patterns through bounded false positives. This approach introduces no tunable parameters, as the canonical partitions are fixed by construction.

In Table VI, the bandwidth overhead—measured as the average ratio of padded to true response size—quantifies the

cost introduced by each defense. Padding to powers-of-two incurs a modest overhead of less than $2\times$, while REMIN shows only minor degradation, maintaining around 70% tolerant match rate with low MSE. For multiple-of- x padding, as x increases both REMIN and the baseline [1] lose accuracy, but the communication overhead also grows substantially (up to nearly $3\times$). In contrast, BT-style canonical obfuscation almost nullifies reconstruction (tolerant match $< 2\%$), yet at a heavy cost: each response requires on average more than $4\times$ the original communication.

Notably, under extreme padding schemes (maximum multiples of x) or BT-style obfuscation, the baseline method essentially fails, whereas REMIN still preserves meaningful reconstruction. This robustness arises from REMIN’s use of co-occurrence statistics: bogus padded records appear far less frequently than true records, and contribute negligible signal, allowing our model to naturally filter them out. In contrast, the baseline approach treats all records equally, making it vulnerable to spurious correlations introduced by padding.

Overall, these results illustrate a clear trade-off: lightweight padding offers limited protection with modest overhead, while strong obfuscation provides higher privacy at prohibitive cost. This aligns with practical SSE deployments, which often prioritize efficiency and remain susceptible to leakage-abuse attacks—underscoring the continued real-world relevance of our threat model. While extremely sensitive information could justify such overhead, in scenarios where access and search patterns cannot be observed, all current attacks targeting these channels are largely ineffective.

VI. REMIN-P ATTACK: LEVERAGING AUXILIARY INFORMATION FOR IMPROVED RECONSTRUCTION

In this section, we explore a new approach for reducing reconstruction error by utilizing auxiliary information. We consider a natural extension of the adversarial setting in which the attacker, beyond passive observation, can poison a limited number of records into the dataset. The technique—our poisoning strategy—operates on the output of any reconstruction algorithm and is agnostic to the underlying inference mechanism. As such, it represents a general augmentation layer that can enhance or exploit the structure recovered by any range query leakage-abuse attack. We describe the new attack, called REMIN-P, and demonstrate its effectiveness in improving the database reconstruction.

A. Intention

We explore how to further minimize our reconstruction error, especially in scenarios where limited query data leads to imperfect reconstructions. In practical settings, some records in the dataset may be known to the attacker, either due to unavoidable leakage or intentional injections. Such scenarios are realistic: an adversary may generate controlled transactions in online services, inject identifiable files into target servers, or—under permissive contexts—introduce or monitor test patients and visits [61], [62], [63], [64], thereby inserting poisoned records into the system. One may use this

TABLE VI: Comparison of defense mechanisms against REMIN and the baseline from Markatou et al. [1] (all experiments under 30×30 grid). Bandwidth overhead is measured as the average ratio of padded to true result size.

Defense	Bandwidth overhead	REMIN (ours)				Markatou et al. [1]			
		MSE \downarrow	Match \uparrow	Neighbor \uparrow	Chamfer \downarrow	MSE \downarrow	Match \uparrow	Neighbor \uparrow	Chamfer \downarrow
Original (no defense)	1.00x	1.49	0.80	0.82	0.48	31.28	0.28	0.64	0.62
Max padding (power-of-2)	1.43x	2.17	0.68	0.77	0.47	28.46	0.27	0.64	0.66
Max padding (multiply-of-50)	1.30x	2.75	0.61	0.80	0.52	148.1	0.01	0.02	1.13
Max padding (multiply-of-75)	1.47x	5.64	0.57	0.75	0.54	149.8	0.01	0.02	0.89
Max padding (multiply-of-125)	1.83x	4.11	0.51	0.76	0.56	155.0	0.00	0.02	1.07
Max padding (multiply-of-250)	2.94x	4.14	0.50	0.76	0.54	271.7	0.01	0.01	2.25
Obfuscation (BT-Style)	4.11x	112.1	0.02	0.46	0.94	214.6	0.01	0.41	2.83

knowledge, referred to as poisoning anchor points, to refine the reconstruction.

We show that even small amounts of auxiliary knowledge can significantly improve the reconstruction of the underlying data. This section explores how poisoned anchor points can be strategically used to adjust the geometry of the reconstructed dataset and reduce edge distortions, with a focus on the impact of two different injection strategies.

B. The REMIN-P Attack Framework

The core idea behind the REMIN-P Attack is to inject auxiliary anchor points into the reconstructed dataset, using known or leaked record positions as spatial references. These anchor points are then used to perform a post-processing adjustment to realign the reconstructed points, improving the overall structure and minimizing distortion (see Algorithm 3).

Algorithm 3 REMIN-P(\hat{P} , Anchors)

Input: \hat{P} : initial reconstruction from REMIN

Anchors: set of poisoned anchor points with known true locations

Output: \hat{P}' : refined reconstruction

```

/* Step 1: Align reconstruction using poisoned anchors */
1:  $\hat{P}_a \leftarrow \text{PROCRUSTESALIGN}(\hat{P}, \text{Anchors})$ 
/* Step 2: Estimate correction function from anchor discrepancies */
2:  $\mathcal{C} \leftarrow \text{ESTIMATEANCHORCORRECTION}(\hat{P}_a, \text{Anchors})$ 
/* Step 3: Apply anchor-based correction to all points */
3:  $\hat{P}' \leftarrow \mathcal{C}(\hat{P}_a)$ 
4: return  $\hat{P}'$ 

```

Such auxiliary information may originate from different sources. We distinguish between two typical threat models:

- **Passive Attacker:** gains access to a small, random subset of ground-truth record positions, either through metadata leakage, user-side exposure, or cross-database matching.
- **Active Attacker:** deliberately injects or identifies structurally meaningful points (e.g., central locations, synthetic users) with known coordinates.

These scenarios motivate two corresponding injection strategies-random anchor selection and cross-shaped anchor

placement (Fig. 8). The former explores the effect of anchor quantity under realistic, uncontrolled leakage, while the latter demonstrates how even a few strategically placed anchors can serve as structural scaffolds for global alignment.

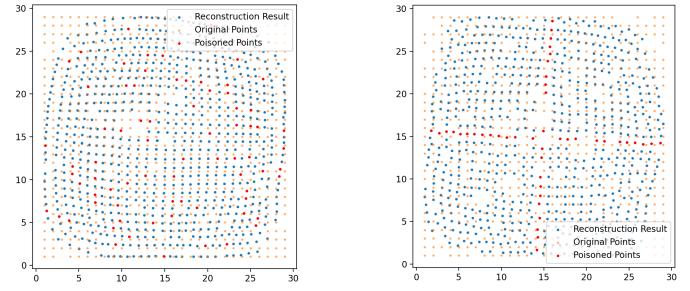


Fig. 8: Comparison of two anchor injection strategies. The left figure illustrates the random placement of poisoned points, while the right shows the cross-shaped method, highlighting the structural improvement achieved by targeting central axes.

Given that the reconstructed embedding may exhibit mild geometric distortions due to manifold learning effects, we propose a post-processing adjustment step that leverages these known anchor points. Specifically, the attacker uses the identified correspondences between a small set of original and reconstructed coordinates to fit a smooth warping function that aligns the entire structure more closely with its true layout. The full procedure is outlined in the full version [51].

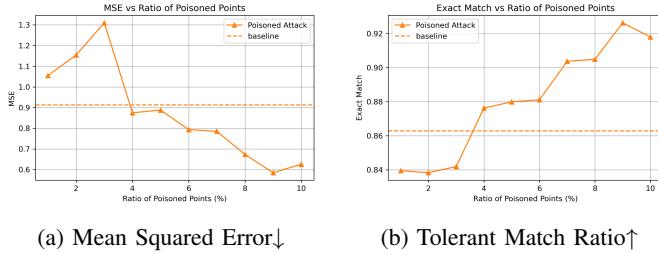
Since edge points typically have fewer neighboring points in the query responses, the positional information extracted for these points tends to be sparser, resulting in greater distortion in the reconstruction edges. To address this, we aim to align the reconstructed space with the true space based on the radial deviation of poisoned points relative to the layout center. By applying 1D interpolation along each axes, the algorithm efficiently warps the reconstructed layout to better match the ground-truth structure. Unlike global affine transformations, this method adapts to local distortions while preserving the relative neighborhood consistency of the reconstruction.

C. Experimental Evaluation

To evaluate the effectiveness of poisoning anchor strategies, we conduct experiments on structured grid datasets under fixed

10% sparsity ranging from 20×20 to 35×35 in size with 5% query ratio sampled uniformly. We simulate the attacker having access to the true coordinates of a small subset of records (anchors), which are used to fit a transformation that adjusts the reconstructed layout.

For the random injection strategy, we fix the grid size (30×30) and vary the anchor point injection ratio across values 1% to 10%. For each ratio, we randomly sample the corresponding number of poisoned points from the dataset. For the cross-shaped injection, we vary the grid size from 20×20 to 40×40 , and insert anchor points along the central row and column (i.e., forming a cross), with the number of poisoned points growing linearly with grid size.



(a) Mean Squared Error \downarrow

(b) Tolerant Match Ratio \uparrow

Fig. 9: Impact of the random injection strategy on reconstruction quality: (Left) MSE and (Right) tolerant match ratio as the ratio of poisoned points increases. Results are obtained on a 30×30 grid dataset with a 5% uniform query ratio.

Our results reveal complementary performance trends under the two anchor injection strategies, with the cross-shaped method achieving strong alignment with fewer anchors, and the random method showing gradual improvement as the injection ratio increases.

Random Injection. Fig. 9 illustrates the quantitative performance of the random poisoning strategy. When the poisoned point ratio is very low (1–3%), the adjustment often introduces additional distortion, degrading the reconstruction accuracy. This is attributed to biased or under-constrained interpolation, where the limited number of anchor points leads to unreliable global correction. However, once the poisoned ratio exceeds 3%, the adjustment consistently outperforms the original reconstruction. As the number of anchors increases, the quality of the correction steadily improves across all evaluation metrics. However, the improvement is gradual, and a large number of anchors are required to reach the performance level achievable by a well-placed set of structural anchors.

Cross-Shaped Injection. In contrast, the cross-shape strategy demonstrates that even limited auxiliary knowledge can be leveraged to significantly improve reconstruction fidelity. As shown in Fig. 10, even with only 19 anchors in a 30×30 grid (lower than 2.5%), we observe a reduction of over 22% in MSE and a gain of 4% in tolerant match accuracy compared to the no-anchor baseline. This targeted selection yields significant structural improvement, particularly in correcting edge distortions. The boundary regions, which often

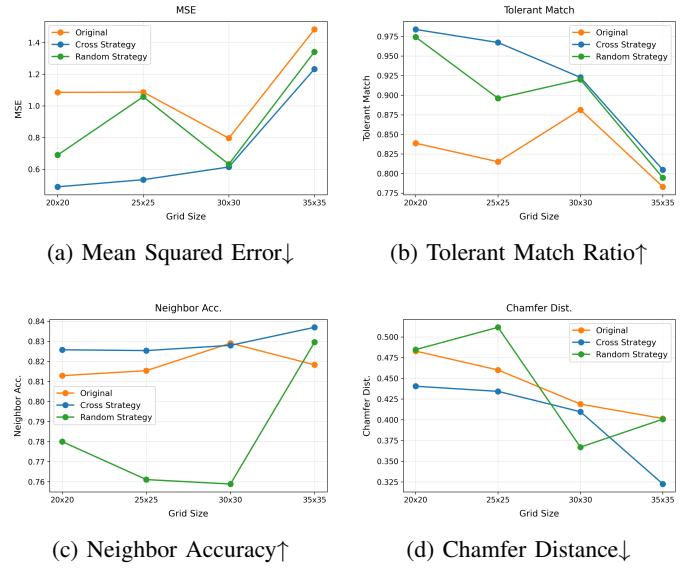


Fig. 10: Comparison of reconstruction metrics before and after post-processing under cross-shaped method. Grid sizes range from 20×20 to 35×35 under 5% uniform query ratio.

exhibit curvature or compression artifacts are realigned more accurately when poisoned points span both axes.

Moreover, as shown in Fig. 10, this method exhibits greater stability and effectiveness across dataset sizes (from 20×20 to 35×35) compared to random poisoning. Notably, while both strategies lead to improvements, the cross-shaped anchors consistently outperform random anchors at low budget levels, especially for correcting global misalignment.

Our experiments confirm that even limited auxiliary knowledge—either leaked or injected—can be leveraged to significantly improve reconstruction fidelity. In practice, it suggests that attackers may benefit from strategically introducing or identifying semantically informative points, thereby refining the geometry of otherwise noisy embeddings. This calls for stronger defense mechanisms beyond traditional leakage profiling, including robustness against minimal anchoring and geometry-aware leakage mitigation.

VII. DISCUSSION AND FUTURE WORK

Challenges in Extremely Sparse Datasets. Experimental observations indicate a dramatic decline in reconstruction performance in highly sparse datasets, characterized by over 70% empty regions. In such scenarios, the reduced co-occurrence of distant records severely limits the attacker’s observable information. This sparsity also increases the likelihood of queries returning very few or even single records, further weakening the statistical signal needed for reconstruction. Future work could explore incorporating additional signals, such as the frequency of individual point occurrences, to estimate the size of empty neighborhoods, thereby improving reconstruction capabilities in sparse data environments.

Framework Limitations and Generalization. While REMIN shows strong performance under extremely limited

query leakage, its scalability under higher leakage ratios (e.g., 50%) remains underexplored. Evaluating whether full database reconstruction is achievable in such settings may further reveal the method’s robustness and practical potential. Moreover, the use of t-SNE introduces sensitivity to parameters such as perplexity and stochastic variability, which may affect reconstruction fidelity—especially in non-uniform datasets where isolated points offer limited relational signals. Although our poisoning-based REMIN-P helps mitigate some of these issues, future work could explore embedding algorithms that incorporate structural constraints or density-aware regularization. More broadly, the modular structure of our framework suggests potential for extending to other types of leakage, by transforming diverse observable signals into learnable spatial representations.

Opportunities and Challenges in Query Poisoning. Our REMIN-P shows that even limited manipulation—through observing or injecting a small number of ground-truth records—can substantially improve reconstruction outcomes. A slightly more active attacker could further exploit this by crafting specific queries or strategically inserting points to guide the embedding process. Future work may explore adaptive poisoning strategies that leverage dataset geometry or embedding sensitivities.

Toward Dynamic Environments. Our analysis focuses a practical setting where records are stored and retrieved in cloud servers, with underlying real-world databases remaining largely stable during evaluation, as commonly assumed in prior SSE attacks [1], [19], [17]. In dynamic environments, insertions and deletions continuously reshape dataset boundaries and range coverage, altering access patterns, response volumes, and co-occurrence relationships. Such temporal drift compounds search–update traces and renders the leakage distribution non-stationary, requiring adaptive calibration instead of a fixed benchmark snapshot [65]. Addressing this non-stationarity remains an open challenge for future work.

VIII. CONCLUSION

In this work, We address the challenge of reconstructing spatial datasets under severe query leakage constraints, a critical problem in privacy-preserving systems and adversarial settings. We present a novel machine learning-based attack that leverages representation learning from highly limited range query leakage. Our method, based on t-SNE and co-occurrence analysis, is the first to demonstrate effective reconstruction under general datasets, outperforming state-of-the-art baselines by up to 50% in MSE. These results highlight a critical vulnerability in current privacy-preserving systems and show that even minimal query exposure can lead to meaningful data recovery, calling for stronger defenses in real-world deployments.

IX. ETHICS CONSIDERATIONS

This work demonstrates the vulnerability of searchable encryption systems to practical database reconstruction attacks using minimal leakage, which can help inspire stronger

security measures and privacy-preserving techniques. By exposing these risks, our research contributes positively to the advancement of secure encrypted database systems. Below we discuss key ethical considerations regarding intellectual property, intended usage, potential misuse, risk control, and human subjects.

Intellectual property. All comparative attacks and defenses, models, datasets and implementation libraries are open-source.

Intended Usage. Our work demonstrates how minimal leakage from encrypted databases can enable approximate data reconstruction, exposing critical vulnerabilities in current systems. These findings aim to drive the development of more secure encryption schemes that better protect sensitive information.

Potential Misuse. Our findings could be misused to reconstruct sensitive location data, infer private financial/medical information, or manipulate data via poisoning attacks. Potential defenses may include padding [66], ORAM [4], [5], and poisoning detection system.

Risk Control. To further mitigate potential risks, we will release the code of this work including algorithms and the implementations in the experiments. We believe that transparency can reduce the risks related to this work, encourage reliable code reuse and promote advancement of database security.

Human Subject. This research does not involve human subjects and any personally identifiable information.

REFERENCES

- [1] E. A. Markatou and R. Tamassia, “Reconstructing with even less: Amplifying leakage and drawing graphs,” in *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, 2024, pp. 4777–4791.
- [2] D. X. Song, D. Wagner, and A. Perrig, “Practical techniques for searches on encrypted data,” in *Proceeding 2000 IEEE symposium on security and privacy. S&P 2000*. IEEE, 2000, pp. 44–55.
- [3] R. Curtmola, J. Garay, S. Kamara, and R. Ostrovsky, “Searchable symmetric encryption: improved definitions and efficient constructions,” in *Proceedings of the 13th ACM conference on Computer and communications security*, 2006, pp. 79–88.
- [4] O. Goldreich and R. Ostrovsky, “Software protection and simulation on oblivious rams,” *Journal of the ACM (JACM)*, vol. 43, no. 3, pp. 431–473, 1996.
- [5] E. Stefanov, M. v. Dijk, E. Shi, T.-H. H. Chan, C. Fletcher, L. Ren, X. Yu, and S. Devadas, “Path oram: an extremely simple oblivious ram protocol,” *Journal of the ACM (JACM)*, vol. 65, no. 4, pp. 1–26, 2018.
- [6] C. Gentry, “Fully homomorphic encryption using ideal lattices,” in *Proceedings of the forty-first annual ACM symposium on Theory of computing*, 2009, pp. 169–178.
- [7] A. Boldyreva, N. Chenette, Y. Lee, and A. O’neill, “Order-preserving symmetric encryption,” in *Advances in Cryptology-EUROCRYPT 2009: 28th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cologne, Germany, April 26-30, 2009. Proceedings 28*. Springer, 2009, pp. 224–241.
- [8] MongoDB, Inc., “Queryable encryption: Technical guide,” MongoDB, Tech. Rep., 2022, accessed: 2025-04-20. [Online]. Available: <https://www.mongodb.com/docs/manual/core/queryable-encryption/>
- [9] J. Yao, Y. Zheng, Y. Guo, and C. Wang, “Sok: A systematic study of attacks in efficient encrypted cloud data search,” in *Proceedings of the 8th International Workshop on Security in Blockchain and Cloud Computing*, 2020, pp. 14–20.
- [10] C. Liu, L. Zhu, M. Wang, and Y.-a. Tan, “Search pattern leakage in searchable encryption: Attacks and new construction,” *Information Sciences*, vol. 265, pp. 176–188, 2014.

[11] M. S. Islam, M. Kuzu, and M. Kantacioglu, “Access pattern disclosure on searchable encryption: ramification, attack and mitigation.” in *Ndss*, vol. 20. Citeseer, 2012, p. 12.

[12] G. Kellaris, G. Kollios, K. Nissim, and A. O’neill, “Generic attacks on secure outsourced databases,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016, pp. 1329–1340.

[13] M.-S. Lacharité, B. Minaud, and K. G. Paterson, “Improved reconstruction attacks on encrypted data using range query leakage,” in *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2018, pp. 297–314.

[14] P. Grubbs, M.-S. Lacharité, B. Minaud, and K. G. Paterson, “Learning to reconstruct: Statistical learning theory and encrypted database attacks,” in *2019 IEEE symposium on security and privacy (SP)*. IEEE, 2019, pp. 1067–1083.

[15] E. M. Kornaropoulos, C. Papamanthou, and R. Tamassia, “The state of the uniform: Attacks on encrypted databases beyond the uniform query distribution,” in *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2020, pp. 1223–1240.

[16] F. Falzon, E. A. Markatou, Akshima, D. Cash, A. Rivkin, J. Stern, and R. Tamassia, “Full database reconstruction in two dimensions,” in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, 2020, pp. 443–460.

[17] E. A. Markatou, F. Falzon, R. Tamassia, and W. Schor, “Reconstructing with less: Leakage abuse attacks in two dimensions,” in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, 2021, pp. 2243–2261.

[18] P. Grubbs, M.-S. Lacharité, B. Minaud, and K. G. Paterson, “Pump up the volume: Practical database reconstruction from volume leakage on range queries,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, 2018, pp. 315–331.

[19] E. A. Markatou and R. Tamassia, “Full database reconstruction with access and search pattern leakage,” in *International Conference on Information Security*. Springer, 2019, pp. 25–43.

[20] E. A. Markatou, F. Falzon, Z. Espiritu, and R. Tamassia, “Attacks on encrypted response-hiding range search schemes in multiple dimensions,” *Proceedings on Privacy Enhancing Technologies*, 2023.

[21] L. Van der Maaten and G. Hinton, “Visualizing data using t-sne.” *Journal of machine learning research*, vol. 9, no. 11, 2008.

[22] R. Bost, “oφoc: Forward secure searchable encryption,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016, pp. 1143–1154.

[23] R. Bost, B. Minaud, and O. Ohrimenko, “Forward and backward private searchable encryption from constrained cryptographic primitives,” in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 2017, pp. 1465–1482.

[24] D. Cash, J. Jaeger, S. Jarecki, C. Jutla, H. Krawczyk, M.-C. Roşu, and M. Steiner, “Dynamic searchable encryption in very-large databases: Data structures and implementation,” *Cryptology ePrint Archive*, 2014.

[25] D. Cash, S. Jarecki, C. Jutla, H. Krawczyk, M.-C. Roşu, and M. Steiner, “Highly-scalable searchable symmetric encryption with support for boolean queries,” in *Advances in Cryptology—CRYPTO 2013: 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2013. Proceedings, Part I*. Springer, 2013, pp. 353–373.

[26] M. Chase and S. Kamara, “Structured encryption and controlled disclosure,” in *International conference on the theory and application of cryptology and information security*. Springer, 2010, pp. 577–594.

[27] M. George, S. Kamara, and T. Moataz, “Structured encryption and dynamic leakage suppression,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2021, pp. 370–396.

[28] J. Ghareh Chamani, D. Papadopoulos, C. Papamanthou, and R. Jalili, “New constructions for forward and backward private symmetric searchable encryption,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, 2018, pp. 1038–1055.

[29] S. Kamara and C. Papamanthou, “Parallel and dynamic searchable symmetric encryption,” in *Financial Cryptography and Data Security: 17th International Conference, FC 2013, Okinawa, Japan, April 1–5, 2013, Revised Selected Papers 17*. Springer, 2013, pp. 258–274.

[30] S. Kamara, C. Papamanthou, and T. Roeder, “Dynamic searchable symmetric encryption,” in *Proceedings of the 2012 ACM conference on Computer and communications security*, 2012, pp. 965–976.

[31] M. Naveed, M. Prabhakaran, and C. A. Gunter, “Dynamic searchable encryption via blind storage,” in *2014 IEEE Symposium on Security and Privacy*. IEEE, 2014, pp. 639–654.

[32] E.-J. Goh, “Secure indexes,” *Cryptology ePrint Archive*, 2003.

[33] P. Golle, J. Staddon, and B. Waters, “Secure conjunctive keyword search over encrypted data,” in *International conference on applied cryptography and network security*. Springer, 2004, pp. 31–45.

[34] L. Ballard, S. Kamara, and F. Monroe, “Achieving efficient conjunctive keyword searches over encrypted data,” in *International conference on information and communications security*. Springer, 2005, pp. 414–426.

[35] Y. H. Hwang and P. J. Lee, “Public key encryption with conjunctive keyword search and its extension to a multi-user system,” in *International conference on pairing-based cryptography*. Springer, 2007, pp. 2–22.

[36] D. Boneh and B. Waters, “Conjunctive, subset, and range queries on encrypted data,” in *Theory of Cryptography: 4th Theory of Cryptography Conference, TCC 2007, Amsterdam, The Netherlands, February 21–24, 2007. Proceedings 4*. Springer, 2007, pp. 535–554.

[37] D. Cash, P. Grubbs, J. Perry, and T. Ristenpart, “Leakage-abuse attacks against searchable encryption,” in *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, 2015, pp. 668–679.

[38] L. Blackstone, S. Kamara, and T. Moataz, “Revisiting leakage abuse attacks,” *Cryptology ePrint Archive*, 2019.

[39] E. M. Kornaropoulos, N. Moyer, C. Papamanthou, and A. Psomas, “Leakage inversion: Towards quantifying privacy in searchable encryption,” in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 1829–1842.

[40] G. Wang, C. Liu, Y. Dong, K.-K. R. Choo, P. Han, H. Pan, and B. Fang, “Leakage models and inference attacks on searchable encryption for cyber-physical social systems,” *IEEE Access*, vol. 6, pp. 21 828–21 839, 2018.

[41] D. Pouliot and C. V. Wright, “The shadow nemesis: Inference attacks on efficiently deployable, efficiently searchable encryption,” in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 1341–1352.

[42] Y. Varabéi, I. Kabin, Z. Dyka, D. Klann, and P. Langendoerfer, “Intelligent clustering as a means to improve k-means based horizontal attacks,” in *2019 IEEE 30th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC Workshops)*. IEEE, 2019, pp. 1–6.

[43] J. Cook, J. Drean, J. Behrens, and M. Yan, “There’s always a bigger fish: A clarifying analysis of a machine-learning-assisted side-channel attack,” in *Proceedings of the 49th Annual International Symposium on Computer Architecture*, 2022, pp. 204–217.

[44] M. A. Carreira-Perpinán, “A review of dimension reduction techniques,” *Department of Computer Science. University of Sheffield. Tech. Rep. CS-96-09*, vol. 9, pp. 1–69, 1997.

[45] L. Van Der Maaten, E. O. Postma, H. J. Van Den Herik *et al.*, “Dimensionality reduction: A comparative review,” *Journal of machine learning research*, vol. 10, no. 66–71, p. 13, 2009.

[46] P. Cunningham, “Dimension reduction,” in *Machine learning techniques for multimedia: Case studies on organization and retrieval*. Springer, 2008, pp. 91–112.

[47] H. Abdi and L. J. Williams, “Principal component analysis,” *Wiley interdisciplinary reviews: computational statistics*, vol. 2, no. 4, pp. 433–459, 2010.

[48] W. S. Torgerson, “Multidimensional scaling: I. theory and method,” *Psychometrika*, vol. 17, no. 4, pp. 401–419, 1952.

[49] G.-X. Tian, M.-J. Cheng, and S.-Y. Cui, “The generalized reciprocal distance matrix of graphs,” *arXiv preprint arXiv:2204.03787*, 2022.

[50] T. Hofmann, B. Schölkopf, and A. J. Smola, “Kernel methods in machine learning,” 2008.

[51] P. Li, H. Chen, E. A. Markatou, and K. Liang, “Learning from leakage: Database reconstruction from just a few multidimensional range queries,” *Cryptology ePrint Archive*, Paper 2025/2248, 2025. [Online]. Available: <https://eprint.iacr.org/2025/2248>

[52] “Estimation of obesity levels based on eating habits and physical condition,” UCI Machine Learning Repository, 2019, DOI: <https://doi.org/10.24432/C5H31Z>.

[53] “Iranian Churn,” UCI Machine Learning Repository, 2020, DOI: <https://doi.org/10.24432/CSJW3Z>.

[54] A. Srinivasan, “Statlog (Landsat Satellite),” UCI Machine Learning Repository, 1993, DOI: <https://doi.org/10.24432/C55887>.

[55] V. Vo, X. Yuan, S.-F. Sun, J. K. Liu, S. Nepal, and C. Wang, “Shielddb: An encrypted document database with padding countermeasures,” *IEEE*

Transactions on Knowledge and Data Engineering, vol. 35, no. 4, pp. 4236–4252, 2021.

[56] S. Han, V. Chakraborty, M. T. Goodrich, S. Mehrotra, and S. Sharma, “Veil: A storage and communication efficient volume-hiding algorithm,” *Proceedings of the ACM on Management of Data*, vol. 1, no. 4, pp. 1–27, 2023.

[57] R. A. Popa, C. M. Redfield, N. Zeldovich, and H. Balakrishnan, “Cryptdb: Protecting confidentiality with encrypted query processing,” in *Proceedings of the twenty-third ACM symposium on operating systems principles*, 2011, pp. 85–100.

[58] I. Demertzis, D. Papadopoulos, C. Papamanthou, and S. Shintre, “Seal: Attack mitigation for encrypted databases via adjustable leakage,” in *29th USENIX security symposium (USENIX Security 20)*, 2020, pp. 2433–2450.

[59] H. Nie, W. Wang, P. Xu, X. Zhang, L. T. Yang, and K. Liang, “Query recovery from easy to hard: Jigsaw attack against {SSE},” in *33rd USENIX Security Symposium (USENIX Security 24)*, 2024, pp. 2599–2616.

[60] G. Chen, T.-H. Lai, M. K. Reiter, and Y. Zhang, “Differentially private access patterns for searchable symmetric encryption,” in *IEEE INFOCOM 2018-IEEE conference on computer communications*. IEEE, 2018, pp. 810–818.

[61] Y. Zhang, J. Katz, and C. Papamanthou, “All your queries are belong to us: the power of {File-Injection} attacks on searchable encryption,” in *25th USENIX Security Symposium (USENIX Security 16)*, 2016, pp. 707–720.

[62] Z. Wang, J. Yu, M. Gao, W. Yuan, G. Ye, S. Sadiq, and H. Yin, “Poisoning attacks and defenses in recommender systems: A survey,” *arXiv preprint arXiv:2406.01022*, 2024.

[63] X. Zhang, W. Wang, P. Xu, L. T. Yang, and K. Liang, “High recovery with fewer injections: Practical binary volumetric injection attacks against dynamic searchable encryption,” in *32nd USENIX Security Symposium (USENIX Security 23)*, 2023, pp. 5953–5970.

[64] T. J. Langhout, H. Chen, and K. Liang, “File-injection attacks on searchable encryption, based on binomial structures,” in *European Symposium on Research in Computer Security*. Springer, 2024, pp. 424–443.

[65] L. Xu, L. Zheng, C. Xu, X. Yuan, and C. Wang, “Leakage-abuse attacks against forward and backward private searchable symmetric encryption,” in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, 2023, pp. 3003–3017.

[66] S. Vaudenay, “Security flaws induced by cbc padding—applications to ssl, ipsec, wtls...,” in *International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2002, pp. 534–545.

APPENDIX A ADDITIONAL MATERIALS

A. Notation Summary

TABLE VII: Notation Summary

Symbol	Description
n	Total number of records
\mathcal{F}	Original dataset with n records
S_d	A d -dimensional Euclidean space
N_i	Range size in dimension i ($i = 1, \dots, d$)
\mathcal{E}	Encrypted dataset stored on the server
t_q	Query token corresponding to the query q
$\text{Response}(t_q)$	Set of returned records of the query token t_q
\mathcal{Q}	Set of observed queries
F	The co-occurrence frequency matrix
$D(i, j)$	The (i, j) -th entry of the distance matrix D
Y	Representation matrix after dimension reduction

B. Supplement Datasets and Performance under Non-uniform Query Distribution

The non-geospatial datasets we evaluate are:

- **Estimation of Obesity Levels:** Records from individuals in Mexico, Peru, and Colombia, used to estimate obesity levels based on lifestyle and physical condition. For 2D evaluation, we use the attributes *height* and *weight*; for 3D evaluation, we additionally include a dietary habit indicator (“vegetable consumption”). We refer to this dataset as *Grid*. *We refer to this dataset as Obesity*.
- **Iranian Churn:** Customer data from an Iranian telecom provider, containing 13 behavioral and service-related features (e.g., usage duration, call failures, subscription length). For 2D evaluation, we use *subscription length* and *seconds of use*; in 3D, we further include *call failure rate*. *We refer to this dataset as Telecom*.
- **Landsat Satellite:** Multi-spectral measurements from satellite imagery, each record representing a 3×3 pixel neighborhood labeled by soil or crop type. In the 2D setting, we use the *grey soil* and *red soil* attributes; in 3D, we add the *cotton crop* dimension. *We refer to this dataset as Landsat*.

All features are normalized and scaled to a 50×50 grid.

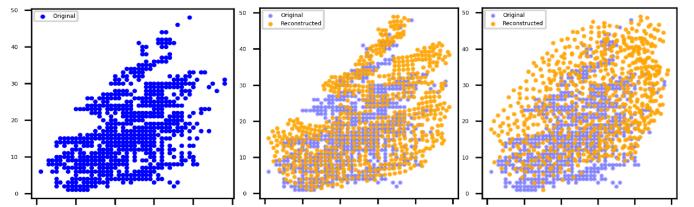


Fig. 11: Visualization of the Obesity dataset (height–weight pairs) scaled to a 50×50 grid. (Middle) Reconstruction by REMIN using 1% uniformly sampled queries, achieving 83.71% neighbor accuracy and MSE of 37.41. (Right) Reconstruction by the method of Markatou et al. [1] under the same setting, achieving 47.30% neighbor accuracy and MSE of 130.07.

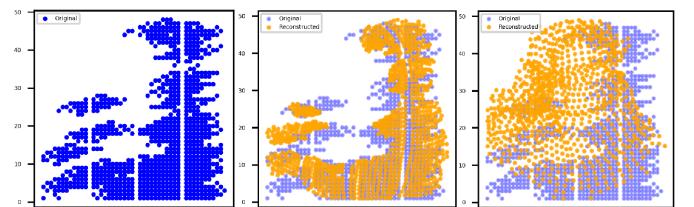


Fig. 12: Visualization of the Iranian Telecom dataset (Subscription Length–Seconds of Use pairs) scaled to a 50×50 grid. (Middle) Reconstruction by REMIN using 1% uniformly sampled queries, achieving 84.46% neighbor accuracy and MSE of 12.89. (Right) Reconstruction by the method of Markatou et al. [1] under the same setting, achieving 45.27% neighbor accuracy and MSE of 200.25.

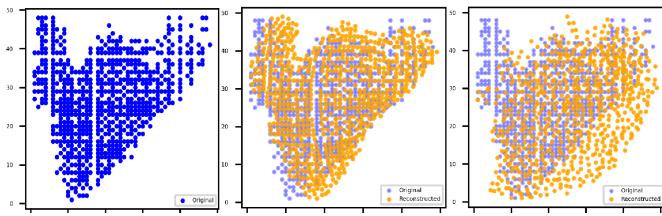


Fig. 13: Visualization of the Landsat Satellite dataset (red soil–grey soil pairs) scaled to a 50×50 grid. (Middle) Reconstruction by REMIN using 1% uniformly sampled queries, achieving 85.51% neighbor accuracy and MSE of 14.78. (Right) Reconstruction by the method of Markatou et al. [1] under the same setting, achieving 54.90% neighbor accuracy and MSE of 63.47.

Unlike spatial datasets, these datasets consist of behavioral and transactional attributes rather than physical coordinates, but still exhibit structured distributions that make them suitable for reconstruction analysis. Under 1% uniform query coverage, the reconstruction results (Fig. 11, 12 and 13) show that our method consistently outperforms the baseline [1], confirming that its effectiveness generalizes across heterogeneous datasets.

To better capture practical query behaviors, we follow SOTA [1], [17], [20] in adopting non-uniform query models, e.g., Beta, Gaussian distributions, which emulate the biased or correlated access patterns commonly observed. Notably, the impact of query distribution lies not in introducing bias but in determining the number of unique queries observed. A more skewed distribution effectively reduces diversity of query observations, thereby limiting leakage signal. As illustrated in Table VIII, our method naturally outperforms the baseline [1], as it is explicitly designed to extract maximal information from co-occurrence frequencies even under extremely limited query observations.

C. Supplementary Experimental Results

To complement the evaluation presented in the main text, we include additional figures in this appendix. These visualizations provide further evidence of the effectiveness and efficiency of our method across different settings.

TABLE IX: Reconstruction MSE under different sparsity levels and dimensions.

Sparsity (%)	2D (64)		3D (16)		4D (8)	
	REMIN	Markatou et al. [1]	REMIN	Markatou et al. [1]	REMIN	Markatou et al. [1]
0	0.2159	-	0.3667	0.7008	0.6499	0.8305
10	0.2946	-	0.4640	1.0090	0.6134	0.6985
20	0.4147	-	0.4323	0.6783	0.6673	0.8712
30	0.5909	-	0.4903	0.8221	0.7426	0.8176
40	0.6283	-	0.6649	1.0821	0.7260	0.9611
50	0.7911	-	0.6843	1.3349	0.8835	0.9747

TABLE VIII: MSE Reduction (%) achieved by our method compared to the method of Markatou et al. [1] under three query distribution settings: Beta, Gaussian, and Uniform.

Distribution	MSE Ours	MSE Markatou et al. [1]	Reduction %
AMS Dataset			
Uniform	23.52	37.86	37.87%
Beta	54.46	112.87	51.75%
Gaussian	62.81	111.58	43.71%
Paris Dataset			
Uniform	42.17	53.07	20.53%
Beta	48.58	120.30	59.61%
Gaussian	57.85	238.37	75.73%
Manhattan Dataset			
Uniform	14.25	42.97	66.83%
Beta	42.87	97.72	56.13%
Gaussian	48.23	157.27	69.34%
Shanghai Dataset			
Uniform	22.81	36.92	38.19%
Beta	38.28	89.72	57.32%
Gaussian	46.55	103.69	55.12%
Obesity Dataset			
Uniform	37.41	130.07	71.23%
Beta	56.02	162.84	65.60%
Gaussian	64.17	135.54	52.65%
Telecom Dataset			
Uniform	12.90	200.25	93.55%
Beta	24.51	174.90	85.99%
Gaussian	42.73	243.96	82.50%
Landsat Dataset			
Uniform	14.78	63.48	76.72%
Beta	16.71	92.28	81.89%
Gaussian	33.05	81.22	59.31%

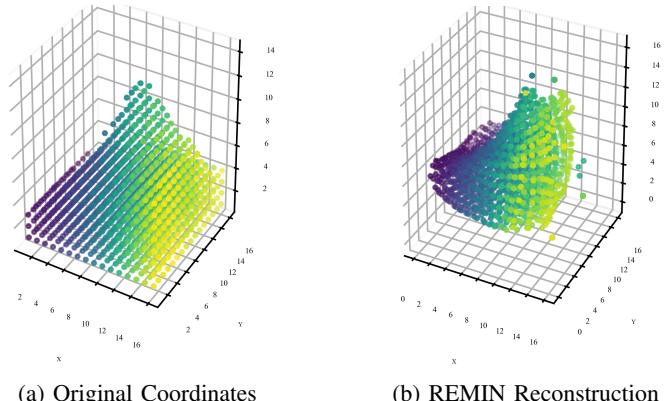


Fig. 14: (Left) Original NH dataset in 3D coordinates. (Right) Reconstruction by our method REMIN using 1% uniformly sampled queries, achieving 72.38% neighbor accuracy, Chamfer distance of 0.5674, and MSE of 0.6745.

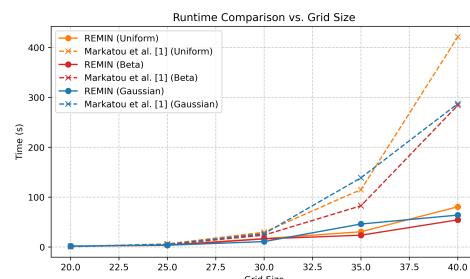


Fig. 15: Runtime across varying grid sizes (20×20 to 40×40) and data distributions with 1% query ratio.

APPENDIX B ARTIFACT APPENDIX

A. Description & Requirements

This artifact provides the complete implementation and experimental environment used to evaluate the REMIN attack. It contains the core reconstruction algorithms, baseline implementations, experiment scripts, and all datasets used in the paper. Evaluators can easily reproduce key results by specifying the dataset, query coverage, and query distribution (uniform, beta, or Gaussian). The pipeline automatically generates query responses, samples observations, and computes reconstruction metrics such as mean squared error (MSE).

1) *How to access*: The artifact is publicly available at: <https://github.com/ZIMUQIN-L/REMIN-attack>. A DOI-enabled archive of the artifact is also available at: <https://doi.org/10.5281/zenodo.17725706>. The repository contains detailed documentation, installation instructions, and example configurations.

2) *Hardware dependencies*: None.

3) *Software dependencies*: The artifact requires a Python 3.8 runtime environment. All dependencies can be installed automatically by running: `pip install -r requirements.txt`. No additional software or system libraries are required beyond standard Python packages.

4) *Benchmarks*: All benchmarks and datasets used in the paper are already included in the artifact package. Evaluators can reproduce any experiment by selecting the desired dataset name and leakage configuration in the provided script parameters (e.g., `-points grid`). No external data or benchmark download is required.

B. Artifact Installation & Configuration

To install and configure the environment for evaluating this artifact, follow the steps below:

- 1) **Obtain the artifact.** Clone or download the artifact repository from the following link: <https://github.com/ZIMUQIN-L/REMIN-attack>
- 2) **Set up the runtime environment.** Ensure that Python 3.8 (or a compatible version) is available on the system. The artifact has been tested on Linux and macOS environments, but it can also run on Windows systems with a standard Python installation.
- 3) **Install dependencies.** From the root directory of the repository, run the following command to automatically install all required packages:

```
pip install -r requirements.txt
```

- 4) **Verify the setup.** Once installation completes successfully, the environment is ready for use. You can confirm the setup by executing a minimal example:

```
python eval_exp.py
```

This test runs a small-scale reconstruction using a toy dataset and verifies that all dependencies are correctly configured.

No additional configuration or system-level setup is required beyond the steps above. The entire artifact can be executed within a standard Python 3.8 environment on a commodity desktop machine.

C. Experiment Workflow

The artifact provides a unified interface to reproduce all experiments described in the paper. The workflow consists of three stages:

- 1) **Step 1: Fetch Response Data.** Generates or loads query-response pairs for a chosen dataset. Configurable parameters include:
 - `points`: select which dataset to use (e.g., `grid`, `drink`, `credit`).
 - `p`: the proportion of observable queries (e.g., 1 for 1% leakage).
 - `dist`: the query distribution type, chosen from `uniform`, `beta`, or `gaussian`.
 - `N0`, `N1`: define the query domain size in each dimension.
- 2) **Step 2: REMIN Reconstruction Attack.** Executes the REMIN or REMIN-P algorithm to reconstruct the dataset from sampled query-response pairs. Optional hyperparameter:
 - `perplexity`: controls t-SNE dimensionality reduction.
 All other parameters use the default values reported in the paper.
- 3) **Step 3: Evaluation.** Automatically compares reconstructed and ground-truth datasets, reporting quantitative metrics (MSE, Chamfer distance, neighborhood accuracy). No configuration is required for this step.

All three stages can be executed end-to-end using a single command:

```
python eval_exp.py -points grid <dataset>
-p <query ratio> -dist <distribution>
-N0 <domain size> -N1 <domain size>
```

This command runs the full workflow—data preparation, reconstruction, and evaluation—without additional manual setup.

D. Major Claims

- (C1): REMIN consistently outperforms the state-of-the-art baseline on both synthetic and real-world datasets, achieving up to 50% lower MSE. Proven by Experiments (E1)–(E2) whose results are reported in Fig. 1, Fig. 3 and Fig. 4.
- (C2): REMIN remains robust under varying sparsity, dimensionality, and query distributions, demonstrating strong scalability and generalization. Proven by Experiments (E1) whose results are reported in Fig. 6, Fig. 7 and Table II.

E. Evaluation

1) *Experiment (E1)—Real-World Datasets:* [20 human-minutes + 3 compute-hour] Validate reconstruction performance on practical datasets under uniform query coverage.

[Preparation] Install dependencies using `pip install -r requirements.txt`.

[Execution] Run both methods with 1% query coverage under uniform query distribution:

```
python eval_exp.py -points drinking -p 1  
-dist uniform -N0 50 -N1 50
```

[Results] REMIN consistently achieves lower MSE and better structural fidelity than the baseline across all datasets, confirming its robustness to real-world data distributions.

2) *Experiment (E2)—Synthetic Datasets:* [1 human-minutes + 4 compute-hour] Evaluate reconstruction accuracy and robustness under controlled sparsity, dimensionality, and query distributions.

[Preparation] Configure the parameters for generating synthetic grid-structured datasets, including grid dimensions (N_0 , N_1) ranging from 20 to 40, and sparsity levels from 0% to 70%.

[Execution] Run both methods under identical query budgets (1%) with query distributions `uniform`, `beta`, and `gaussian`:

```
python eval_exp.py -points grid -p 1  
-dist uniform -N0 20 -N1 20
```

Repeat each configuration with 3–5 seeds to account for t-SNE randomness.

[Results] Expected outcome: REMIN achieves consistently lower MSE than the baseline (up to 50% reduction), with stable performance across sparsity and query distributions. Accuracy remains above 50% even with 60% sparsity, confirming scalability and generalization.