Privacy Preserving Identification Using Sparse Approximation with Ambiguization

Behrooz Razeghi, Slava Voloshynovskiy, Dimche Kostadinov and Olga Taran

Stochastic Information Processing Group
University of Geneva
Switzerland

December 2017

UNIVERSITÉ DE GENÈVE

SIP stochastic information processing

Outline

Introduction

Privacy-preserving content identification

- Biometrics
- Physical object recognition and security
- Medical/clinical applications
- Privacy-sensitive multimedia records

Introduction

Privacy-preserving content identification

- Biometrics
- Physical object recognition and security
- Medical/clinical applications
- Privacy-sensitive multimedia records

Recent Trends

Big Data & Distributed Applications

Services on outsourced cloud-based systems

Introduction

Privacy-preserving content identification

- Biometrics
- Physical object recognition and security
- Medical/clinical applications
- Privacy-sensitive multimedia records

Recent Trends

Big Data & Distributed Applications

Services on outsourced cloud-based systems

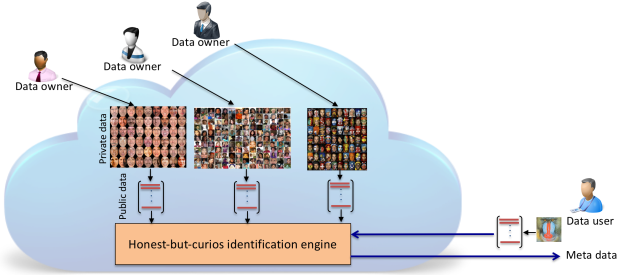

Introduction

Problem Formulation

Goal of privacy protection in outsourced services

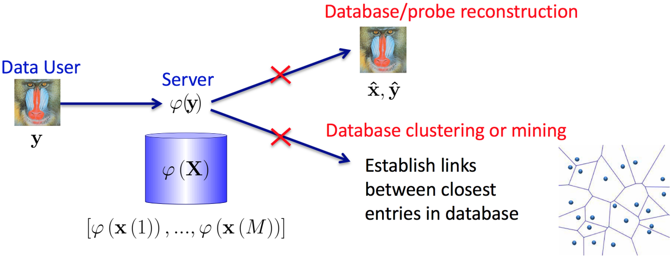

Introduction

How do we receive a feature vector?

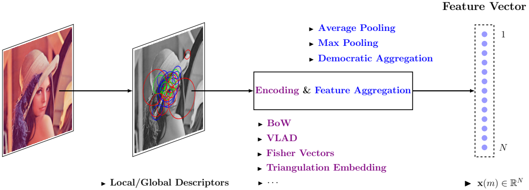

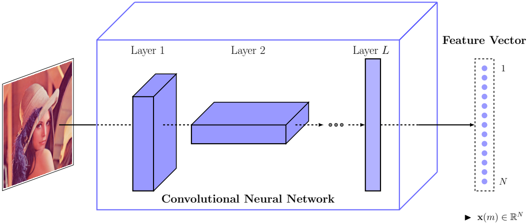

Introduction

state-of-the-art

- **Cryptographic Methods - Homomorphic Encryption**
  - **Main Idea:** Similarity search in the encrypted domain
    - Brute force identification $\Longrightarrow$ huge complexity
- **Robust Hashing - a single hash from the whole content / local descriptors / last layer of CNN**
  - **Main Idea:** $\mathbf{x} \longrightarrow (011011100110)$ and believed non-invertability
    - Loss in performance due to binarization
    - Unauthorized database clustering

Introduction

state-of-the-art

- **Cryptographic Methods - Homomorphic Encryption**
  - **Main Idea:** Similarity search in the encrypted domain
    - Brute force identification $\Longrightarrow$ huge complexity
- **Robust Hashing - a single hash from the whole content / local descriptors / last layer of CNN**
  - **Main Idea:** $\mathbf{x} \longrightarrow (011011100110)$ and believed non-invertability
    - Loss in performance due to binarization
    - Unauthorized database clustering
- **Group Testing / Memory Vectors**
  - **Main Idea:** Group testing by measuring the proximity to the group representative
    - Group representatives (memory vectors) should be stored in memory

Introduction

state-of-the-art

- **Cryptographic Methods - Homomorphic Encryption**
  - **Main Idea:** Similarity search in the encrypted domain
    - Brute force identification $\Longrightarrow$ huge complexity
- **Robust Hashing - a single hash from the whole content / local descriptors / last layer of CNN**
  - **Main Idea:** $\mathbf{x} \longrightarrow (011011100110)$ and believed non-invertability
    - Loss in performance due to binarization
    - Unauthorized database clustering
- **Group Testing / Memory Vectors**
  - **Main Idea:** Group testing by measuring the proximity to the group representative
    - Group representatives (memory vectors) should be stored in memory

Introduction

state-of-the-art

Universal Quantization

- **Main Idea:** projection with the dimension reduction and periodic quantization
  - Binary quantization: in the region of low projected magnitudes - high P_b
  - Ambiguization due to periodization of quantizer - no possibility to recover data even for the authorized users
  - Server can still can cluster data - privacy leakages
  - Information preservation in general - no link to $R(d)$ and recovery is demonstrated so far

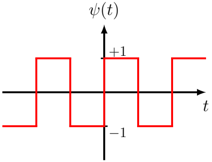

$$\mathbf{a} = \psi\left(\mathbf{W}\mathbf{x}\right)$$

$$t_i = \left[\mathbf{W}\mathbf{x}\right]_i$$

Introduction

state-of-the-art

- **Proposed approach: 3 key elements**
  - Sparsification
  - Ambiguization
  - Search / Identification
- **Advantages:**
  - Fast search / memory efficient
  - Difficult to accurately reconstruct from probe
  - Server cannot reveal a structure of the database

Introduction

state-of-the-art

- **Proposed approach: 3 key elements**
  - Sparsification
  - Ambiguization
  - Search / Identification
- **Advantages:**
  - Fast search / memory efficient
  - Difficult to accurately reconstruct from probe
  - Server cannot reveal a structure of the database
- **Main concerns addressed in our study:**
  - Performance
  - Memory (database) / complexity (identification)
  - Privacy-preserving with respect to:
    - database $\mathcal{A}$
    - probe $\mathbf{y}$

Introduction

state-of-the-art

- **Proposed approach: 3 key elements**
  - Sparsification
  - Ambiguization
  - Search / Identification
- **Advantages:**
  - Fast search / memory efficient
  - Difficult to accurately reconstruct from probe
  - Server cannot reveal a structure of the database
- **Main concerns addressed in our study:**
  - Performance
  - Memory (database) / complexity (identification)
  - Privacy-preserving with respect to:
    - database $\mathcal{A}$
    - probe $\mathbf{y}$

Part 1:
Sparse Data Representation

Sparsification

Main Idea

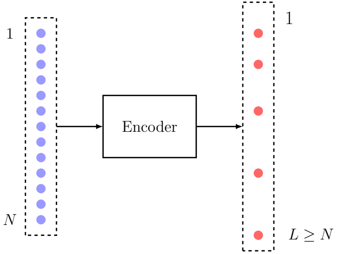

- $\mathbf{x}(m) \in \mathbb{R}^N$
- $\mathbf{x}(m) \sim p(\mathbf{x})$

- $\mathbf{a}(m) \in \{-1, 0, +1\}^L$
- $\|\mathbf{a}(m)\|_0 \leq S_x$
- Rate: $R = \frac{1}{L} \log_2 \left(\binom{L}{S_x} 2^{S_x} \right)$

Sparsification

Main Idea

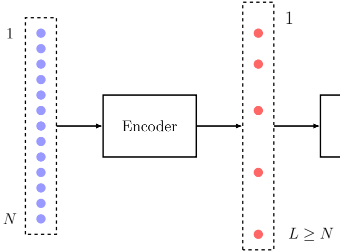

- $\mathbf{x}(m) \in \mathbb{R}^N$
- $\mathbf{x}(m) \sim p(\mathbf{x})$
- $\mathbf{a}(m) \in \{-1, 0, +1\}^L$
- $\|\mathbf{a}(m)\|_0 \leq S_x$
- Rate: $R = \frac{1}{L}\log_2\left(\binom{L}{S_x}2^{S_x}\right)$
- $\hat{\mathbf{x}}(m) \in \mathbb{R}^N$
- Distortion : $\frac{1}{N}\|\mathbf{x}(m) - \hat{\mathbf{x}}(m)\|_2^2 \leq D$

Sparsification

Main Idea

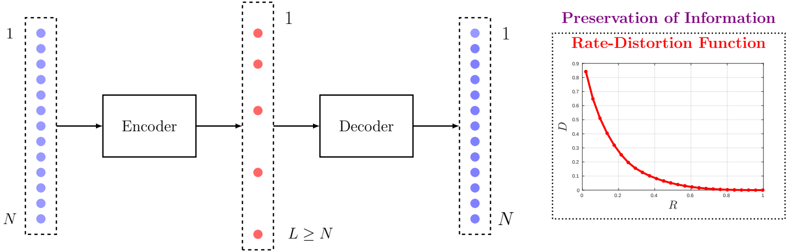

- $\mathbf{x}(m) \in \mathbb{R}^N$
- $\mathbf{x}(m) \sim p(\mathbf{x})$
- $\mathbf{a}(m) \in \{-1, 0, +1\}^L$
- $\|\mathbf{a}(m)\|_0 \leq S_x$
- Rate: $R = \frac{1}{L} \log_2 \left(\binom{L}{S_x} 2^{S_x} \right)$
- $\hat{\mathbf{x}}(m) \in \mathbb{R}^N$
- Distortion : $\frac{1}{N} \|\mathbf{x}(m) - \hat{\mathbf{x}}(m)\|_2^2 \leq D$

Part 2:
Ambiguization

Ambiguization

Main Idea

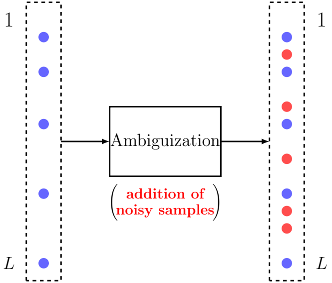

- $\mathbf{a}(m) \in \{-1, 0, +1\}^L$
- $\|\mathbf{a}(m)\|_0 \le S_x$
- **Public Domain**
- $\mathbf{a}(m) \bigoplus \mathbf{n}$

Ambiguization

Main Idea

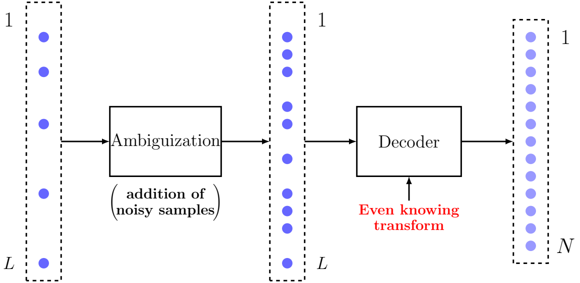

- $\mathbf{a}(m) \in \{-1, 0, +1\}^L$
- $\|\mathbf{a}(m)\|_0 \leq S_x$

- **Public Domain**
- $\mathbf{a}(m) \bigoplus \mathbf{n}$

- $\hat{\mathbf{x}}(m) \in \mathbb{R}^N$
- $\|\mathbf{x}(m) - \hat{\mathbf{x}}(m)\|_2^2 \rightarrow \simeq 2N\sigma_{\mathbf{x}}^2$

- Prevent reconstruction from $\mathbf{a}(m) \bigoplus \mathbf{n}$ and from probe $\mathbf{y}$
- Preclude server from discovering the structure of the database $\mathcal{A}$

Part 3:
Privacy-Preserving Identification

Privacy-Preserving Identification: Private Search

Main Idea: User discloses his probe completely

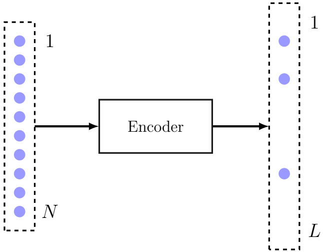

- $\mathbf{y} \in \mathbb{R}^N$
- $\mathbf{b} \in \{-1, 0, +1\}^L$
- $\|\mathbf{b}\|_0 \leq S_y$

Privacy-Preserving Identification: Private Search

Main Idea: User discloses his probe completely

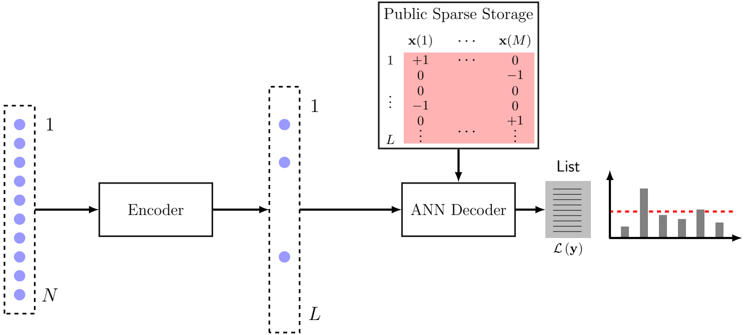

- $\mathbf{y} \in \mathbb{R}^N$
- $\mathbf{b} \in \{-1, 0, +1\}^L$
- $\|\mathbf{b}\|_0 \leq S_y$

Privacy-Preserving Identification: Public Search

Main Idea: User sends only positions of interest

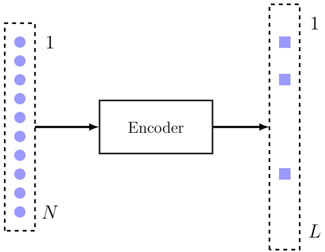

- $\mathbf{y} \in \mathbb{R}^N$
- $\mathbf{b} \in \{-1, 0, +1\}^L$
- $\|\mathbf{b}\|_0 \leq S_y$

Privacy-Preserving Identification: Public Search

Main Idea: User sends only positions of interest

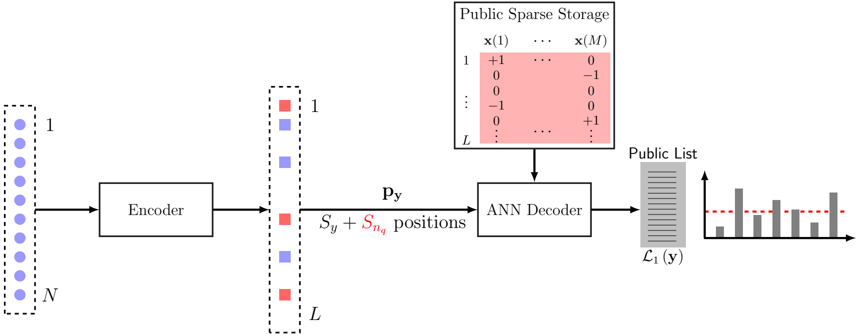

- $\mathbf{y} \in \mathbb{R}^N$
- $\mathbf{b} \in \{-1, 0, +1\}^L$
- $\|\mathbf{b}\|_0 \leq S_y$
- Add S_{n_q} random positions

Privacy-Preserving Identification: Public Search

Main Idea: User sends only positions of interest

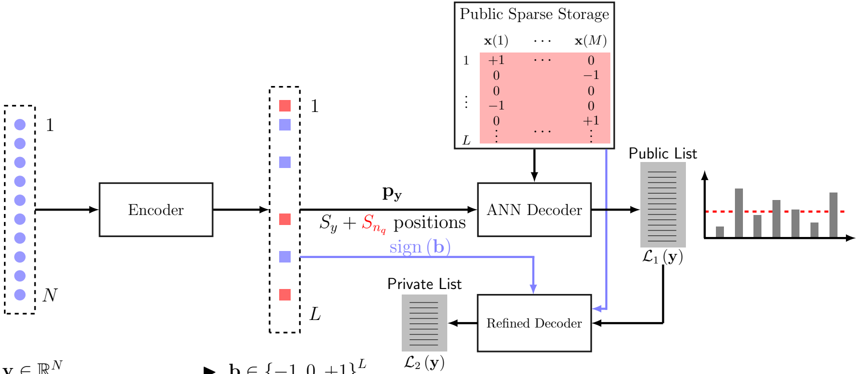

- $\mathbf{y} \in \mathbb{R}^N$
- $\mathbf{b} \in \{-1, 0, +1\}^L$
- $\|\mathbf{b}\|_0 \leq S_y$
- Add S_{n_q} random positions

Main idea behind the proposed solution

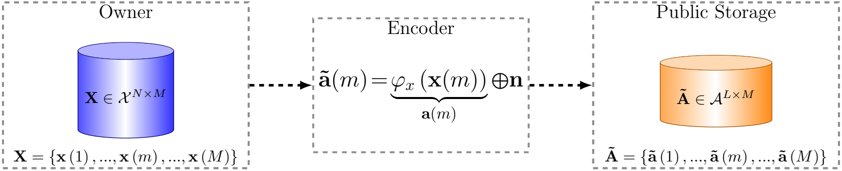

Main idea behind the proposed solution

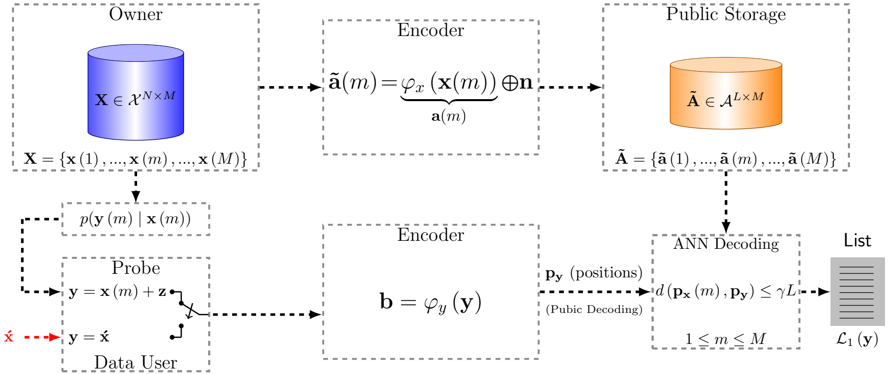

Main idea behind the proposed solution

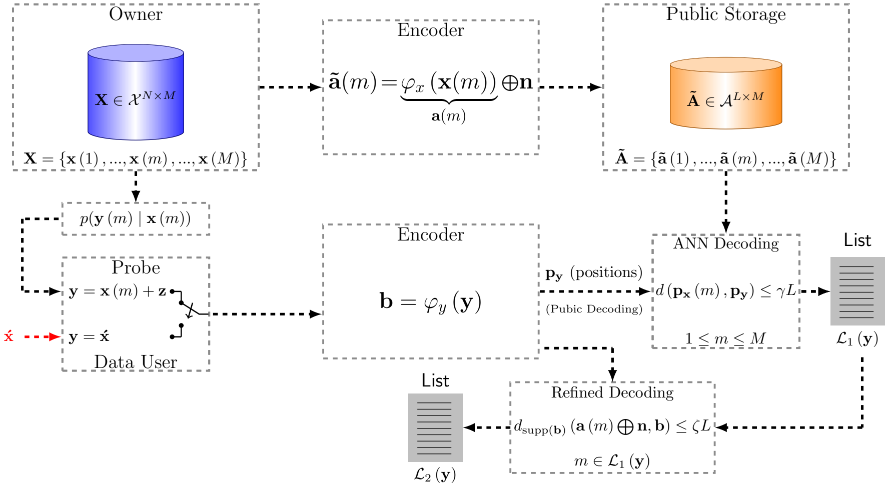

Sparsifying Transform

A Schematic Idea

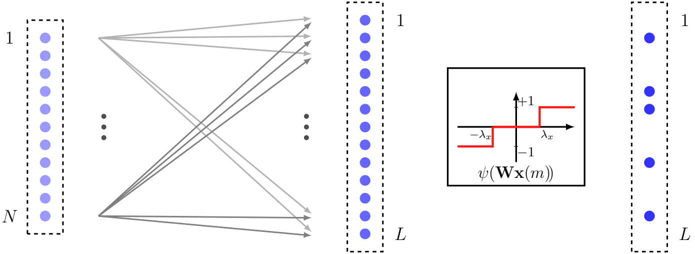

Linear Mapping

Element-wise Non-linearity

$$\begin{tikzcd} \mathbf{x}(m) \in \mathbb{R}^N \arrow[r, "\mathbf{W}"] \arrow[rr, bend right, "\varphi(.)"'] & \mathbf{Wx}(m) \in \mathbb{R}^L \arrow[r, "\psi(.)"] & \mathbf{a}(m) \in \{-1, 0, +1\}^L \end{tikzcd}$$

Sparsifying Transform

General Problem Formulation

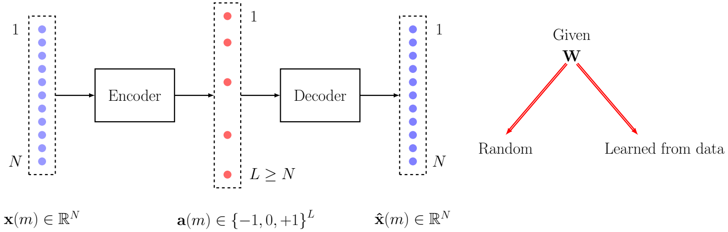

- **Encoder:**

$$\boxed{\hat{\mathbf{a}}\left(m\right) = \psi\left(\mathbf{W}\mathbf{x}\left(m\right)\right)}$$

- **Decoder:**

$$\boxed{\hat{\mathbf{x}}\left(m\right) = \mathbf{W}^{\dagger}\hat{\mathbf{a}}\left(m\right)}$$

Encoder: as a projection problem (for a fixed $\mathbf{W}$)

$$\widehat{\mathbf{a}}(m) = \underset{\mathbf{a}(m)\in\mathcal{A}^L}{\arg\min} \; \|\mathbf{W}\mathbf{x}(m) - \mathbf{a}(m)\|_2^2 + \beta\Omega\left(\mathbf{a}(m)\right), \forall m \in [M]$$

- $\mathbf{W} \in \mathbb{R}^{L\times N}$, $\mathbf{x}(m) \in \mathbb{R}^N$, $\mathbf{a}(m) \in \mathbb{R}^L$
- Closed-form solution for: $\Omega\left(.\right) = \|.\|_0$ and $\Omega\left(.\right) = \|.\|_1$

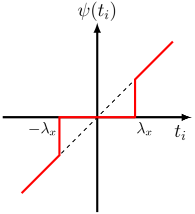

Hard-thresholding operator

$\Omega\left(.\right) = \|.\|_0$

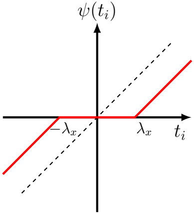

Soft-thresholding operator

$\Omega\left(.\right) = \|.\|_1$

$t_i = \left[\mathbf{W}\mathbf{x}\right]_i$

$$\hat{\mathbf{a}}\left(m\right) = \psi\left(\mathbf{W}\mathbf{x}\left(m\right)\right)$$

Encoder: Extra constraint on the alphabet

$$\boxed{\hat{\mathbf{a}}(m) = \psi(\mathbf{W}\mathbf{x}(m))}$$

$$\text{s.t.} \qquad \mathbf{a}(m) \in \{-1, 0, +1\}$$

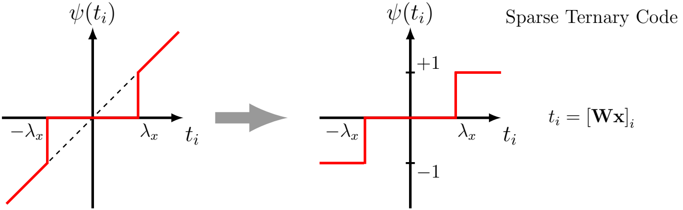

Encoder: Extra constraint on the alphabet

$$\boxed{\hat{\mathbf{a}}\,(m) = \psi\left(\mathbf{W}\mathbf{x}\,(m)\right)}$$

$$\text{s.t.} \qquad \mathbf{a}\,(m) \in \{-1, 0, +1\}$$

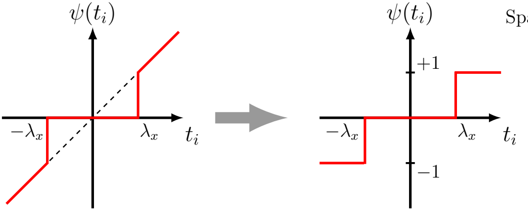

Sparse Ternary Code

$t_i = [\mathbf{W}\mathbf{x}]_i$

Remark:

Binary hashing (like LSH) is the special case of our $\psi(.)$ for $\lambda_x = 0$.

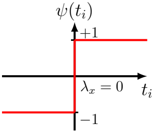

Comparison of Three Encoding Schemes

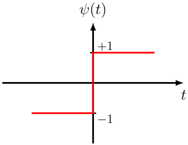

Binary Hashing

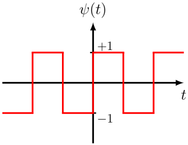

Quantized Embeddings

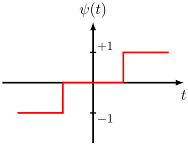

Sparse Ternary Coding

Learning Sparsifying Transform

General Formulation: joint learning

$$\left(\hat{\mathbf{W}}, \hat{\mathbf{A}}\right) = \arg\min_{(\mathbf{W},\mathbf{A})} \|\mathbf{W}\mathbf{X} - \mathbf{A}\|_F^2 + \beta_W \Omega_W(\mathbf{W}) + \beta_A \Omega_A(\mathbf{A})$$

- **Sparse Coding Step** (Fixed $\mathbf{W}$):

$$\hat{\mathbf{A}} = \arg\min_{\mathbf{A}} \|\mathbf{W}\mathbf{X} - \mathbf{A}\|_F^2 + \beta_A \Omega_A(\mathbf{A})$$

$$\boxed{\hat{\mathbf{a}}(m) = \psi(\mathbf{W}\mathbf{x}(m))}$$

- **Transform Update Step** (Fixed $\mathbf{A}$):

$$\hat{\mathbf{W}} = \arg\min_{\mathbf{W}} \|\mathbf{W}\mathbf{X} - \mathbf{A}\|_F^2 + \beta_W \Omega_W(\mathbf{W})$$

Linear Regression :
(with quadratic regularizer)

$$\boxed{\hat{\mathbf{W}} = \mathbf{A}\mathbf{X}^T\left(\mathbf{X}\mathbf{X}^T + \beta_W \mathbf{I}_N\right)^{-1}}$$

Ambiguization Scheme

Main Idea

Add noise to **non-zero** components of sparse representation

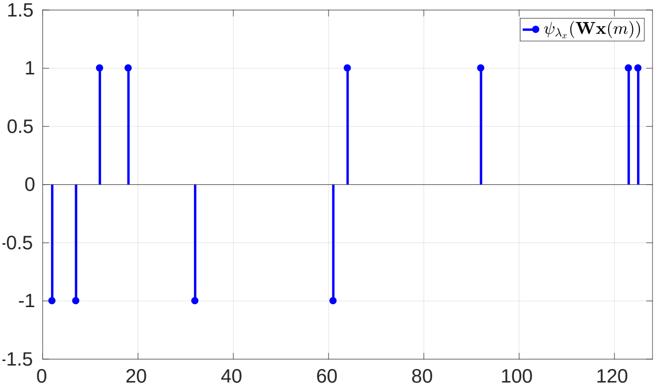

Ambiguization Scheme

Main Idea

Add noise to **non-zero** components of sparse representation

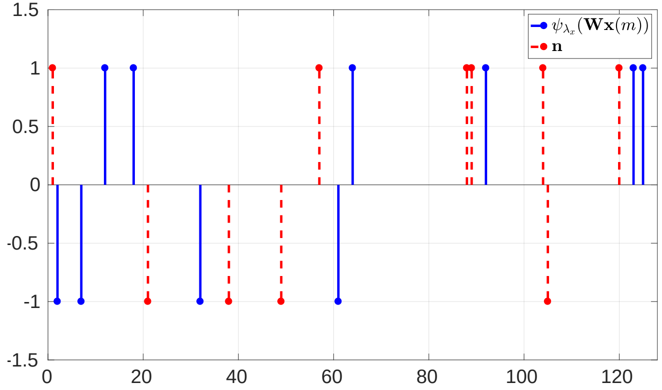

Desired property of mapping scheme

Distance preservation in the desired radius

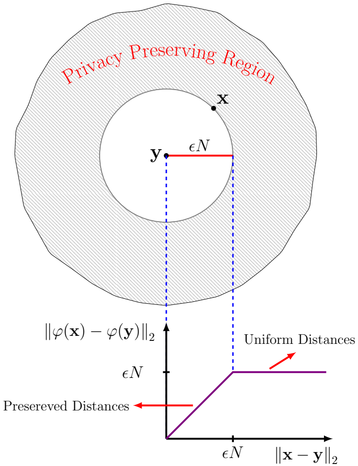

Impact of Ambiguization at Server Side

Goal: The server should not distinguish distances $\| \left(\psi\left(\mathbf{Wx} \right) \bigoplus \mathbf{n} \right) - \psi\left(\mathbf{Wy} \right) \|_2$

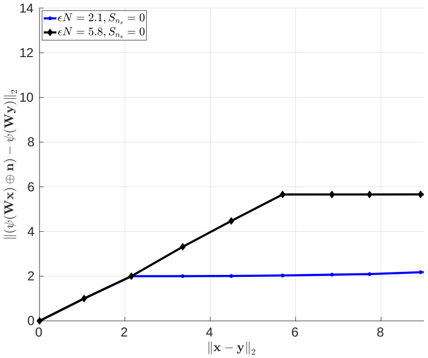

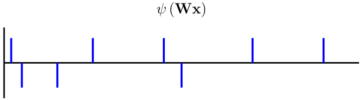

Distances are computed in the full length.

Impact of Ambiguization at Server Side

Goal: The server should not distinguish distances $\| (\psi(\mathbf{Wx}) \oplus \mathbf{n}) - \psi(\mathbf{Wy}) \|_2$

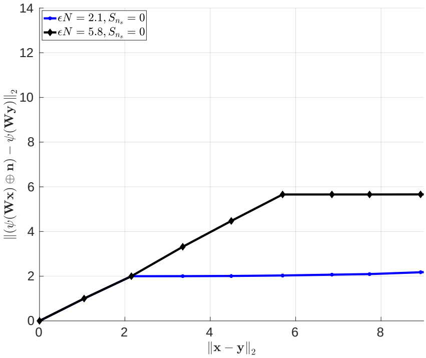

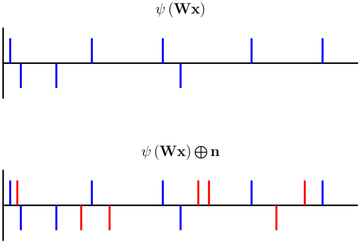

Distances are computed in the full length.

Impact of Ambiguization at Server Side

Goal: The server should not distinguish distances $\| \left(\psi \left(\mathbf{Wx} \right) \bigoplus \mathbf{n} \right) - \psi \left(\mathbf{Wy} \right) \|_2$

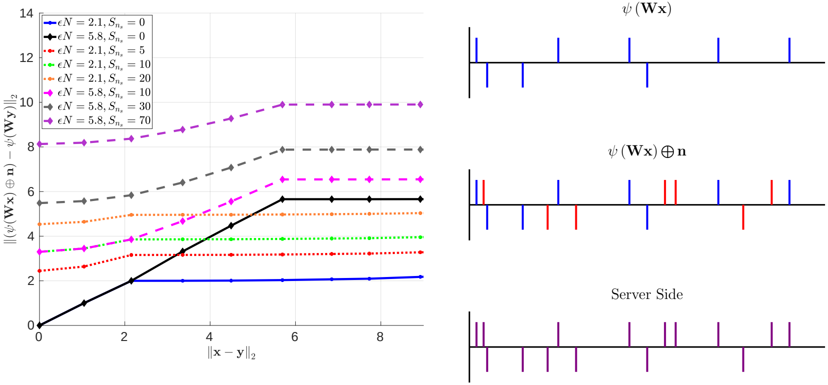

Distances are computed in the full length.

Impact of Ambiguization at Client Side

Goal: The client should distinguish distances $\left(\| \left(\psi\left(\mathbf{Wx}\right) \bigoplus \mathbf{n}\right) - \psi\left(\mathbf{Wy}\right) \|_2\right)_{\text{supp}}$

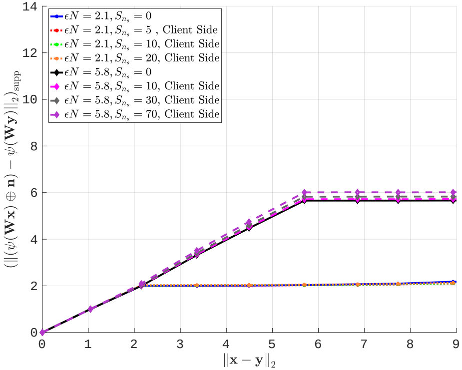

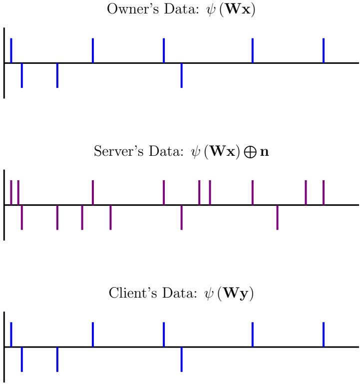

Distances are computed in the non-zero components of probe.

Reconstruction: Authorized User ▶ $\hat{\mathbf{x}} = \mathbf{W}^{\dagger}\mathbf{a}$

$\mathbf{x}$: i.i.d. Gaussian, with each sample $X_n \sim \mathcal{N}(0,1)$, $\frac{N}{L} = 1$

S_x : Sparsity Level

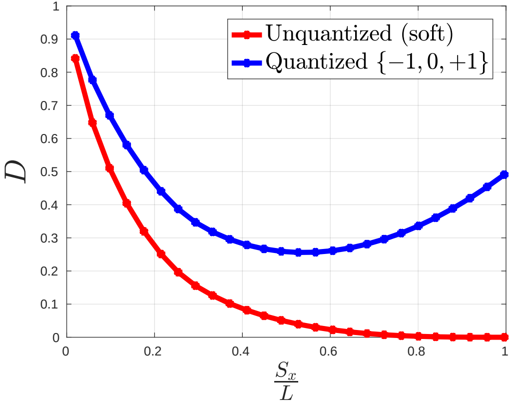

Reconstruction: Unauthorized User ▶ $\hat{\mathbf{x}} = \mathbf{W}^{\dagger}\left(\mathbf{a} \bigoplus \mathbf{n}\right)$

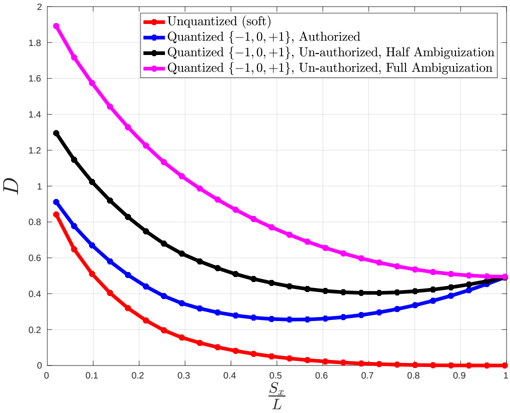

Half Ambiguization: $S_{n_s} = \frac{1}{2}(L - S_x)$ Full Ambiguization: $S_{n_s} = (L - S_x)$

Clustering

Generate Structured Data

- Generate:
  - Four 512-dimensional i.i.d. vectors with distribution $\mathcal{N}\left(\mathbf{0},\mathbf{1}\right)$
  - 1000 512-dimensional i.i.d. vectors with distribution $\mathcal{N}\left(\mathbf{0},\mathbf{0.1}\right)$
- Add each 250 (out of 1000) low variance vectors to the four high variance ones

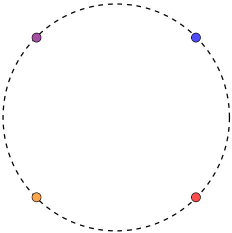

Clustering

Generate Structured Data

- Generate:
  - Four 512-dimensional i.i.d. vectors with distribution $\mathcal{N}\left(\mathbf{0},\mathbf{1}\right)$
  - 1000 512-dimensional i.i.d. vectors with distribution $\mathcal{N}\left(\mathbf{0},\mathbf{0.1}\right)$
- Add each 250 (out of 1000) low variance vectors to the four high variance ones

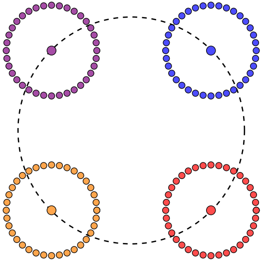

Clustering

Goal: Hide structure of database

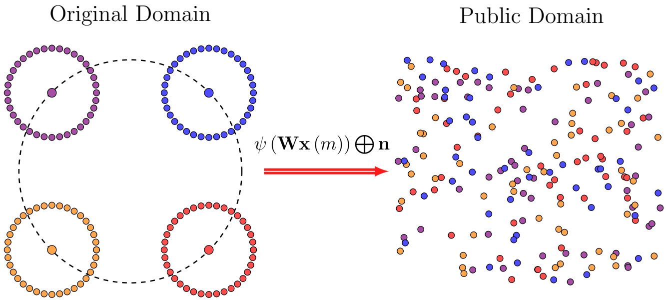

Clustering

Introduce Measure for Evaluation

Define:

- $\alpha_x = \frac{S_x}{L}$, S_x : Sparsity level

Denote:

- P_{intra} : PDF of 'intra-cluster' distances
- P_{inter} : PDF of 'inter-cluster' distances

Define:

- $P_1 = \alpha_x \, P_{\text{intra}} + (1 - \alpha_x) \, P_{\text{inter}}, \;\; 0 \leq \alpha_x \leq 1$

Denote:

- $P_2 \sim \mathcal{N}\left(\mu_2, \sigma_2^2\right)$, fit to P_1

Define:

- Privacy Leak Measure:

$$\begin{aligned} D\left(P_1 \| P_2\right) &= \alpha_x \, D\left(P_{\text{intra}} \| P_2\right) + (1 - \alpha_x) \, D\left(P_{\text{inter}} \| P_2\right) \\ &= \mathbb{E}_{P_1}\left[\log \frac{P_1}{P_2}\right] \end{aligned}$$

Clustering

Introduce Measure for Evaluation

Define:

- $\alpha_x = \frac{S_x}{L}$, S_x : Sparsity level

Denote:

- P_{intra} : PDF of ‘intra-cluster’ distances
- P_{inter} : PDF of ‘inter-cluster’ distances

Define:

- $P_1 = \alpha_x \, P_{\text{intra}} + (1 - \alpha_x) \, P_{\text{inter}}, \quad 0 \leq \alpha_x \leq 1$

Denote:

- $P_2 \sim \mathcal{N}\left(\mu_2, \sigma_2^2\right)$, fit to P_1

Define:

- Privacy Leak Measure:

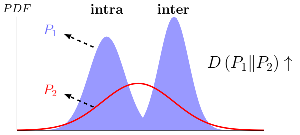

Clear distinguishability based on inter&intra-distances

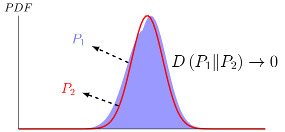

Not distinguishable

$$
\begin{aligned}
D\left(P_1 \| P_2\right) &= \alpha_x \, D\left(P_{\text{intra}} \| P_2\right) + (1 - \alpha_x) \, D\left(P_{\text{inter}} \| P_2\right) \\
&= \mathbb{E}_{P_1}\left[\log \frac{P_1}{P_2}\right]
\end{aligned}
$$

Clustering: How much ambiguization should be added to have indistinguishability for the server?

Evaluation of Our Scheme: $\alpha_x = \frac{S_x}{L}$, $\beta_x = \frac{S_{n_s}}{L}$, S_{n_s} : # of noise components for the server

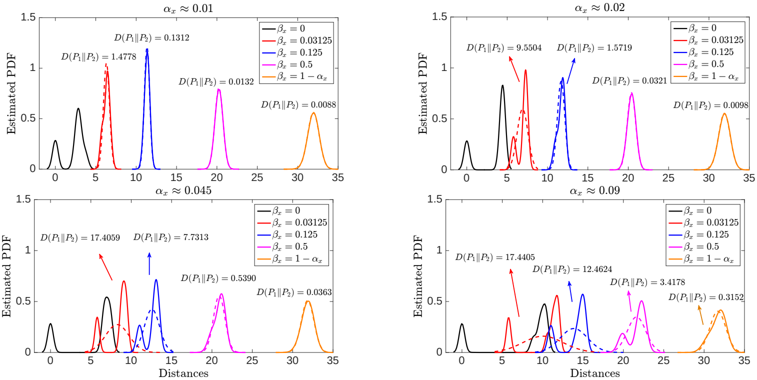

Conclusions:

- Preserve distances up to the desired radius
- Ensure the reconstruction of data for authorized users
- Preclude the curious server to cluster or reconstruct the samples in the database
- Public decoding scheme


Thank You!