

Privacy Preserving Identification Using Sparse Approximation with Ambiguization

Behrooz Razeghi, Slava Voloshynovskiy, Dimche Kostadinov and Olga Taran

Stochastic Information Processing Group, Department of Computer Science, University of Geneva, Switzerland
 {behrooz.razeghi, svolos, dimche.kostadinov, olga.taran}@unige.ch

Abstract—In this paper, we consider a privacy preserving encoding framework for identification applications covering biometrics, physical object security and the Internet of Things (IoT). The proposed framework is based on a sparsifying transform, which consists of a *trained linear map*, an *element-wise nonlinearity*, and *privacy amplification*. The sparsifying transform and privacy amplification are not symmetric for the data owner and data user. We demonstrate that the proposed approach is closely related to sparse ternary codes (STC), a recent information-theoretic concept proposed for fast approximate nearest neighbor (ANN) search in high dimensional feature spaces that being machine learning in nature also offers significant benefits in comparison to sparse approximation and binary embedding approaches. We demonstrate that the privacy of the database outsourced to a server as well as the privacy of the data user are preserved at a low computational cost, storage and communication burdens.

Keywords—*privacy; identification; sparse approximation; transform learning; ambiguization; clustering.*

I. INTRODUCTION

A. Identification and ANN Search

Many modern applications such as biometrics, digital physical object security and data generated by connected objects in the IoT require privacy preserving identification of a query with respect to a given dataset. Practically, the identification problem is based on an ANN search when a list of indices corresponding to the NN items is returned. At the final refinement stage, the list can be refined in a private setting and a single index is declared as the identified one. The identification problem faces the curse of dimensionality. For this reason, the exact identification is replaced by a search of list of closest items, i.e., one tries to trade-off the accuracy of identification by the search complexity. In recent years, many methods providing efficient ANN solutions for multi-billion entry datasets were proposed and we named some of them without pretending to be exhaustive in our overview [1]–[3].

B. Search in Privacy Preserving Settings: Main Considerations

Due to the massive amount of data, modern distributed storage and computing facilities, many ANN problems are considered in a setting where the data user outsources his datasets by applying the corresponding protection measures to third parties (servers) possessing powerful storage, communications and computing facilities. The need for data protection comes from many perspectives related to the cost of data collection, data as a “product” that represents a great value in the era of machine learning, which can be used to train and prune new and existing machine learning tools. Moreover, the server might want to discover some hidden relationships in the data. Finally, the non-renewable nature of some features such as

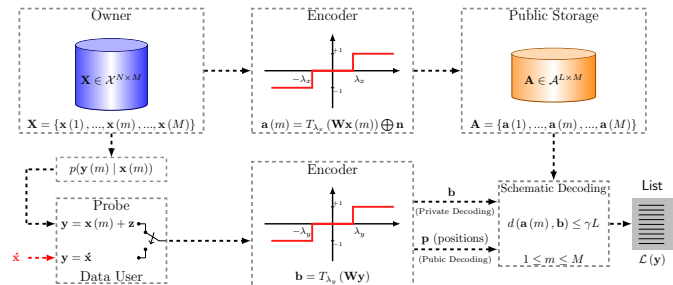


Fig. 1: Block diagram of the proposed model.

for example biometrics, which being disclosed once, do not represent any more a value for the related security applications. The data users (clients), who by hypothesis poses some query, related to those stored in the databases of the data owners, wish to identify them by obtaining an index/indices of the closest items in the data owner datasets. However, the client does not want to disclose his query to the server completely for privacy reasons. Therefore, it is assumed that both the data owner and clients attempt at protecting their data from server side analysis, which is assumed to be honest but curious.

More particularly, the server might attempt: (a) to try to find relationships between entries in the database of the data owner, (b) reconstruct an individual entry of the data owner database or a common representatives or centroids for the clusters, (c) reconstruct the query of clients and (d) cluster multiple queries from the same client or from multiple ones thus establishing the group interests based on the similarity of probes. Curious clients might be also interested to discover more information about the structure of database by exploring the NN around one or multiple probes. Additionally, one can envision collaborative clients who might aggregate the results for multiple queries.

C. Our contribution

This paper presents a new privacy preserving strategy to the ANN search problem based on a recently proposed concept of Sparse Ternary Codes [1], [4]. Our main contribution consists in a novel formulation of database protection based on a STC representation with the addition of ambiguization noise to prevent database analysis and a novel formulation of the query mechanism based on sending the positions of non-zero components in the query sparse representation along with the fraction of ambiguous positions to prevent the reconstruction from query and collaborative processing of queries. Our approach differs from existing privacy protection based on quantized embedding [2], [5] and attribute based [6] techniques. It has strong information-theoretic foundations and demonstrated higher coding gain with respect to binarized embedding methods [4].

Our setup is quite generic and we assume that as an input we might have raw data, extracted features using any known hand crafted methods, aggregated local descriptors based on BoW, FV, VLAD [7]–[9], etc, or intermediate or last layers of deep nets [10]. Before we provide the input to the server we preprocess it by passing it via the sparsifying transform and thresholding it thus producing STC.¹ The next stage is an ambiguation that consists in the addition of selective noise components to the zero positions in the sparse data representation (Fig. 1). As a sparsifying transform we use an efficient learning method based on the sparsified Procrustes problem formulation [11] in contrast to a fixed transform used in [1], [4]. The resulting data are stored on the *public* server. The client, who possesses some query, sends a set of non-zero positions after sparsification of the query to the server and expects that the server will return a set of corresponding indices corresponding to positive and negative non-zero components. To ambiguate the server, the data user also adds a certain portion of false positions to the query and keeps their indices. The search complexity of the server is logarithmic in terms of the number of stored items times the number of non-zero positions in the sparsified query. This also determines the returned file size to the client. Due to the efficient sparse representation, the amount of returned data is minimized while preserving the efficiency of the data structure representation. At the final stage, the data user disregards the lists corresponding to false positions and aggregates the remaining lists for the corresponding positions producing the final ANN list.

D. Notation and Definitions

Throughout this paper, superscript $(\cdot)^T$ stands for the transpose and $(\cdot)^\dagger$ stands for pseudo-inverse. Vectors and matrices are denoted by boldface lower-case (\mathbf{x}) and upper-case (\mathbf{X}) letters, respectively. We consider the same notation for a random vector \mathbf{x} and its realization. The difference should be clear from the context. x_i denotes the i -th entry of vector \mathbf{x} . For a matrix \mathbf{X} , $\mathbf{x}(j)$ denotes the j -th column of \mathbf{X} . Also, we use the notations $[N]$ for the set $\{1, 2, 3, \dots, N\}$ and $\text{card}(\mathcal{S})$ for the cardinality of a set \mathcal{S} .

Definition 1. The support of a vector $\mathbf{x} \in \mathbb{R}^N$ is the index set of its nonzero entries, i.e., $\text{supp}(\mathbf{x}) := \{n \in [N] : x_n \neq 0\}$.

Definition 2. Given a subset \mathcal{E} of a space \mathcal{X} , the indicator function $\mathbb{1}_{\mathcal{E}} : \mathcal{X} \rightarrow \mathbb{R}$ is defined by setting $\mathbb{1}_{\mathcal{E}}(x)$ equal to 1 for $x \in \mathcal{E}$ and equal to 0 for $x \notin \mathcal{E}$.

E. Outline of the Paper

The remainder of the paper is organized as follows. In Section II, the identification problem is defined. We then propose our framework in Section III. Section IV evaluates the privacy imposed by our framework. Finally, conclusions are drawn in Section V.

II. IDENTIFICATION PROBLEM

Consider that M (raw) feature vectors are stored in the database $\mathbf{X} = [\mathbf{x}(1) \ \cdots \ \mathbf{x}(m) \ \cdots \ \mathbf{x}(M)]$, where each feature vector $\mathbf{x}(m)$, $m \in [M]$ from a set $\mathcal{X} \subset \mathbb{R}^N$ is a random vector with distribution $p(\mathbf{x})$ and bounded variance $\sigma_{\mathbf{x}}^2$. The query (probe) $\mathbf{y}(m) \in \mathbb{R}^N$ is the noisy version of $\mathbf{x}(m)$, i.e.,

$\mathbf{y}(m) = \mathbf{x}(m) + \mathbf{z}$, where we assume $\mathbf{z} \in \mathbb{R}^N$ is a Gaussian noise vector with distribution $\mathcal{N}(\mathbf{0}, \Sigma_{\mathbf{z}} = \sigma_{\mathbf{z}}^2 \mathbf{I}_N)$. Provided that the additive noise \mathbf{z} is independent of data, for large N the law of large numbers states that $\frac{1}{N} \|\mathbf{y}(m) - \mathbf{x}(m)\|_2^2 \rightarrow \frac{1}{N} \mathbb{E}[\mathbf{z}^T \mathbf{z}] \simeq \sigma_{\mathbf{z}}^2$, $\forall m \in [M]$. Therefore, for an enrollment vector $\mathbf{x}(m)$ stored at database with high probability query vector $\mathbf{y}(m)$ will be in an N -dimensional sphere of radius $\sqrt{N}\sigma_{\mathbf{z}}$ and centered at $\mathbf{x}(m)$.

III. PROPOSED FRAMEWORK

A. Sparse Representation and Encoding

The intrinsic information content of feature vectors, like other real-world signals, is much smaller than their lengths. So we are interested in approximating them by sparse vectors. Many signals $\mathbf{x}(m) \in \mathbb{R}^N$ may be represented as a linear combination of a small number of columns (words) from a frame² (dictionary) $\mathbf{D} \in \mathbb{R}^{N \times L}$ according to the sparse approximation model. Therefore, $\mathbf{x}(m) = \mathbf{D}\mathbf{a}(m) + \mathbf{e}_{\mathbf{x}}$, where $\mathbf{e}_{\mathbf{x}} \in \mathbb{R}^N$ is approximation error, $\mathbf{a}(m) \in \mathbb{R}^L$ is sparse with $\|\mathbf{a}(m)\|_0 \ll L$, in which $\|\mathbf{a}(m)\|_0 := \text{card}(\text{supp}(\mathbf{a}(m)))$. The general problem formulation of this sparse coding problem for a fixed dictionary can be expressed as:

$$\hat{\mathbf{a}}(m) = \arg \min_{\mathbf{a}(m) \in \mathcal{A}^L} \|\mathbf{x}(m) - \mathbf{D}\mathbf{a}(m)\|_2^2 + \lambda \Omega(\mathbf{a}(m)), \forall m \in [M], \quad (1)$$

where \mathcal{A} is an alphabet of sparse representations, λ is a regularization parameter and $\Omega(\cdot)$ is sparsity-inducing regularization function. However, the sparse coding with respect to $\mathbf{a}(m)$ is an inverse problem in nature and known to be NP-hard. This problem can be solved approximately by greedy or relaxation algorithms, but these only provide the correct solution under certain conditions, which are often violated in real world applications. Furthermore, they are computationally expensive in practice, particularly for large-scale problems.

In this paper, we use a transform model [12] for the sparse representation of vectors at enrollment and identification phases. This model suggests that a feature vector $\mathbf{x}(m) \in \mathbb{R}^N$ is approximately sparsifiable using a transform $\mathbf{W} \in \mathbb{R}^{L \times N}$, that is $\mathbf{W}\mathbf{x}(m) = \mathbf{a}(m) + \mathbf{e}_{\mathbf{a}}$, where $\mathbf{a}(m) \in \mathbb{R}^L$ is sparse, i.e., $\|\mathbf{a}(m)\|_0 \ll L$, and $\mathbf{e}_{\mathbf{a}} \in \mathbb{R}^L$ is the representation error of the feature vector or residual in the *transform domain*. The sparse coding problem for this model is a direct constraint projection problem. The general problem formulation, also known as sparse approximation, is as follows:

$$\hat{\mathbf{a}}(m) = \arg \min_{\mathbf{a}(m) \in \mathcal{A}^L} \|\mathbf{W}\mathbf{x}(m) - \mathbf{a}(m)\|_2^2 + \lambda \Omega(\mathbf{a}(m)), \forall m \in [M]. \quad (2)$$

Note that the residual and therefore the sparse representation $\mathbf{a}(m)$ is not constrained to lie in the range space of \mathbf{W} . Furthermore, since it is a direct problem, for any of the two important regularizers $\Omega(\cdot) = \|\cdot\|_0$ or $\Omega(\cdot) = \|\cdot\|_1$ has closed-form solutions. In this paper, we consider l_0 -“norm” as our sparsity-inducing penalty. In this case, the solution $\hat{\mathbf{a}}(m)$ is obtained exactly by hard-thresholding the projection $\mathbf{W}\mathbf{x}(m)$ and keeping the S_x entries of the largest magnitudes, setting the remaining low magnitude entries to zero. For this purpose, we define an intermediate vector $\mathbf{f}(m) \triangleq \mathbf{W}\mathbf{x}(m) \in \mathbb{R}^L$

¹It should be noted that even higher flexibility is possible for multilevel quantization.

²A set of L orthonormal vectors with vector dimensionality N equal to L is said to form a basis set for that vector space. A frame of an inner product space is a generalization of a basis of a vector space to sets that may be linearly dependent.

and denote by λ_x the S_x -th largest magnitude among the set $\{|f_1(m)|, \dots, |f_L(m)|\}$. Then the closed-form solution is achieved by applying to $\mathbf{f}(m)$ a hard-thresholding operator, which is defined as $H_{\lambda_x}(\mathbf{f}(m)) = \mathbb{1}_{|f_l(m)| \geq \lambda_x} \mathbf{f}(m), \forall m \in [M], \forall l \in [L]$.

In [1], the authors show that Sparse Ternary Coding (STC) outperformed classical binary hashing. Therefore, by considering the alphabet of sparse representation vectors as $\mathcal{A} = \{-1, 0, +1\}$, we apply the ternary hash mapping to $H_{\lambda_x}(\mathbf{W}\mathbf{x}(m))$ as:

$$\mathbf{a}(m) \triangleq T_{\lambda_x}(\mathbf{W}\mathbf{x}(m)) \in \{-1, 0, +1\}^L, \forall m \in [M], \quad (3)$$

where $T_{\lambda_x}(\mathbf{W}\mathbf{x}(m)) = \text{sign}(H_{\lambda_x}(\mathbf{W}\mathbf{x}(m)))$. We denote, therefore, the space of public storage as $\mathcal{A}^{L \times M}$. For simplicity, through out this paper, we denote by $\mathcal{S} \subset \mathcal{X}^N$ the space of vectors in the signal (original) domain and by $\mathcal{T} \subset \mathcal{A}^L$ the space of vectors in the transform domain. Also, we denote by $\psi(\cdot)$ the operator $T_{\lambda}(\cdot)$ in general.

Note that given the transform matrix \mathbf{W} and sparse feature code $\mathbf{a}(m)$, the recovered feature vector is then simply $\hat{\mathbf{x}}(m) = \mathbf{W}^\dagger \mathbf{a}(m)$. We are especially interested in square, i.e., $L = M$, and over-complete, i.e., $L > M$, transform matrices. In this paper, we just address the square sparsifying transform. Obviously, the projection $T_{\lambda_x}(\cdot)$ introduces a certain loss of information. In order to measure the difference between the original data $\mathbf{x}(m)$ and the corresponding reconstructed data $\hat{\mathbf{x}}(m)$ via encoded values $\mathbf{a}(m)$, we use the mean squared error (MSE) as our distortion measure. Also, since our points are random vectors, the distortion measure is given as $D(\mathbf{x}(m), \hat{\mathbf{x}}(m)) = \frac{1}{N} \mathbb{E} [\|\mathbf{x}(m) - \hat{\mathbf{x}}(m)\|_2^2]$. Sparsity level S_x should be chosen such to the specific distortion level.

Our sparsifying transform learning is based on the classical Procrustes matrix problem [13]. That is, we seek an orthonormal matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$, which most closely transforms a given matrix $\mathbf{X} \in \mathbb{R}^{N \times M}$ into a sparse matrix $\mathbf{A} \in \mathbb{R}^{N \times M}$. Therefore, using the Frobenius norm, the problem is to find \mathbf{W} minimizing $\|\mathbf{W}\mathbf{X} - \mathbf{A}\|_F^2$ subject to $\mathbf{W}\mathbf{W}^T = \mathbf{I}$. For the square sparsifying transform, it can be shown that $\mathbf{W} = \mathbf{U}\mathbf{V}^T$, where the columns in \mathbf{U} and \mathbf{V} are left-hand and right-hand singular vectors for $\mathbf{A}\mathbf{X}^T$. Note that, in this case \mathbf{W}^\dagger is simply \mathbf{W}^T . Our algorithm for the above minimization problem alternates between solving for $\mathbf{A} = \psi(\mathbf{W}\mathbf{X})$ (sparse coding step) and $\mathbf{W} = \mathbf{U}\mathbf{V}^T$ (transform update step), whilst the other variables are kept fixed.

In our framework, we set the sparsifying transform matrix \mathbf{W} to be provided, i.e., only authorized parties have access to it. However, even if \mathbf{W} is compromised, our ambiguization mechanism prevents to reconstruct \mathbf{x} from \mathbf{a} . In the general case, the matrix \mathbf{W} is applied selectively to the data and shared only between the concerned parties. The plurality of different matrices $\{\mathbf{W}\}$ is ensured by a key-base initialization of the sparsified Procrustes procedure, i.e., $\mathbf{W}^{(0)} = f(\mathbf{K})$. We choose $f(\mathbf{K}) = \mathbf{U}_K$, where \mathbf{U}_K is unitary matrix obtained from $\mathbf{U}_K \Sigma_K \mathbf{V}_K^T = \mathbf{K}$. In general, given the database $\mathbf{X} \in \mathbb{R}^{N \times M}$ and the shared key matrix $\mathbf{K} \in \mathbb{R}^{L \times N}$ as initial seed for generating \mathbf{W} , the *public sparsifying transform matrix* $\mathbf{W} \in \mathbb{R}^{L \times N}$ can be easily learnt and accessed by authorized parties. Next, given λ_x the encoded ambiguized data can be shared in public storage. Before our discussion on the imposed ambiguization for the server in Sec. IV, note that even if the

server has the clean version of the transformed data, since it has no access to the shared key matrix \mathbf{K} , it cannot reconstruct the original data by the inverse projection. In Sec. IV, we will propose a simple method such that given the disclosed private key \mathbf{K} and public sparsifying transform \mathbf{W} at the server, it is infeasible to learn the structure of data.

B. Link to Quantized Embedding Methods

Many methods like [5] propose a similar privacy preserving framework based on:

$$\mathbf{a}(m) = Q(\mathbf{W}\mathbf{x}(m) + \mathbf{\Delta}), \forall m \in [M], \quad (4)$$

where $Q(\cdot)$ stands for a scalar quantizer and $\mathbf{\Delta}$ is a dither that might be treated as a secret key and $\mathcal{A} = \{-1, +1\}$. There are several fundamental differences between (3) and (4):

- Transform \mathbf{W} : In (4), \mathbf{W} is a dimensionality reduction transform with random entries. In (3), the transform \mathbf{W} is considered as an over-complete transform that extends the dimensions. Moreover, \mathbf{W} in (3) is trained using the described sparsifying Procrustes problem to ensure an optimal sparse representation that is information preserving in general, whereas \mathbf{W} in (4) might preserve the distances only under certain conditions of the Johnson-Lindenstrauss Lemma.
- Codes \mathbf{a} : In (4), the codes are dense and binary, whereas in (3) the codes are sparse and ternary, which form a basis of our ambiguization framework.

C. Identification Stage

At the identification stage, given a query vector $\mathbf{y} \in \mathbb{R}^N$, sparsifying transform matrix \mathbf{W} and sparsity level S_y , vector $\mathbf{b}(m) \triangleq T_{\lambda_y}(\mathbf{W}\mathbf{y}(m)) \in \{-1, 0, +1\}^L$ is calculated. We consider two hypothesis for our query, as follows:

$$\begin{cases} \mathcal{H}_1: & \mathbf{y} = \mathbf{x}(m) + \mathbf{z}, \forall m \in [M], \\ \mathcal{H}_0: & \mathbf{y} = \hat{\mathbf{x}}, \end{cases} \quad (5)$$

where $\hat{\mathbf{x}} \in \mathbb{R}^N$ is a query vector out of the database. The block diagram of our model is depicted in Fig. 1.

D. Practical Decoding

Our final decision is based on the γ -NN decoder, which seeks all $\{\mathbf{a}(m), m \in [M]\}$ NNs in the radius γL from the query \mathbf{y} in order to produce a list of possible candidates as: $\mathcal{L}(\mathbf{y}) = \{m \in [M] : d_{\mathcal{A}}(\mathbf{a}(m), \mathbf{b}) \leq \gamma L\}$, where $d_{\mathcal{A}}(\cdot, \cdot)$ is our similarity measure in space \mathcal{A} . We consider two decoding schemes, *private* and *public* decoding. In the private decoding, the client discloses his probe sparse representation, i.e., $T_{\lambda_y}(\mathbf{W}\mathbf{y}(m))$, to the server completely. While in the public scheme, the client sends only positions of non-zero components in the sparse representation of his probe along with a fraction of ambiguous positions as \mathbf{p} .

E. Isometric Mapping

In general, for a set \mathcal{X} , if we can introduce a metric (or distance function) $d_{\mathcal{X}}$ on $\mathcal{X} \times \mathcal{X}$ which obeying four fundamental axioms [14], then a pair $(\mathcal{X}, d_{\mathcal{X}})$ is said to be a metric space. In our framework we take the set of all ordered N -tuples of \mathcal{S} as our original data, and the set of all ordered L -tuples of \mathcal{T} as our transformed data. Let $\mathcal{S} = (\mathcal{S}, d_{\mathcal{S}})$, $\mathcal{T} = (\mathcal{T}, d_{\mathcal{T}})$ be metric spaces, then a mapping $\psi(\cdot)$ of \mathcal{S} into \mathcal{T} is said to be isometric, if $\psi(\cdot)$ preserves distances.




Fig. 2: The relation between the required sparsity level S_y normalized to code length $L=N$ versus the variance of noise σ_z^2 normalized to σ_x^2 in order to preserve distances in the ternary sparse domain, for $M=20000$ i.i.d random vectors with mean zero and variance σ_x^2 , with three different dimension of feature vectors $N=500, 1000, 1500$ and two sparsity level $S_x=10, 40$.

We are interested to preserve distances by our encoding scheme, i.e., there exists a sufficiently small $\epsilon > 0$ such that for all $m \in [M]$, we have:

$$(1 - \epsilon) d_S(\mathbf{x}(m), \mathbf{y}(m)) \leq d_{\mathcal{T}}(\psi(\mathbf{x}(m)), \psi(\mathbf{y}(m))) \leq (1 + \epsilon) d_S(\mathbf{x}(m), \mathbf{y}(m)). \quad (6)$$

Assuming our metric is induced by the Euclidean norm, we have $d_S(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_2$, $d_{\mathcal{T}}(\mathbf{a}, \mathbf{b}) = \|\mathbf{a} - \mathbf{b}\|_2$. Therefore, noting $\|\mathbf{x}(m) - \mathbf{y}(m)\|_2^2 = \|\mathbf{z}\|_2^2 \simeq N\sigma_z^2$ and considering a ternary sparse encoding, we have:

$$(1 - \epsilon) \|\mathbf{z}\|_2^2 \leq \|\mathbf{a}(m) - \mathbf{b}(m)\|_2^2 \leq (1 + \epsilon) \|\mathbf{z}\|_2^2. \quad (7)$$

Before discussion on the above inequality, we consider lower and upper bounds of distances, which we will use in Section IV. In the original signal domain, considering query vector \mathbf{y} in the general case, the lower and upper bounds of distances between the vector $\mathbf{x}(m)$, $m \in [M]$ in the database and query \mathbf{y} can be given as $N\sigma_z^2 \simeq \|\mathbf{z}\|_2^2 \leq \|\mathbf{x}(m) - \mathbf{y}\|_2^2 \leq \|\mathbf{x}(m)\|_2^2 + \|\mathbf{y}\|_2^2 \simeq N(\sigma_x^2 + \sigma_y^2)$.

In the transform domain with the ternary encoding, we have $\mathbf{W}\mathbf{x}(m) = T_{\lambda_x}(\mathbf{W}\mathbf{x}(m)) + \mathbf{e}_x = \mathbf{a}(m) + \mathbf{e}_x$ and $\mathbf{W}\mathbf{y}(m) = T_{\lambda_y}(\mathbf{W}\mathbf{y}(m)) + \mathbf{e}_y = \mathbf{b}(m) + \mathbf{e}_y$, where \mathbf{e}_x and \mathbf{e}_y are approximation errors of the ternary hash mapping corresponding to $\mathbf{x}(m)$ and $\mathbf{y}(m)$, respectively. Therefore, the upper bound on $\|\mathbf{a}(m) - \mathbf{b}(m)\|_2$ is given as:

$$\|\mathbf{a}(m) - \mathbf{b}(m)\|_2^2 \leq \|\mathbf{W}\|_F^2 \|\mathbf{x}(m) - \mathbf{y}(m)\|_2^2 + \|\mathbf{e}_x - \mathbf{e}_y\|_2^2. \quad (8)$$

However, the above upper bound is not tight, since it does not consider properties of vectors in the transform domain. Also, it is based on the error terms of approximation, which may not give a clear sense of the upper bound at first glance. Considering the sparse ternary property of vectors, we have $\|\mathbf{a}(m) - \mathbf{b}(m)\|_2^2 = \|\mathbf{a}(m)\|_2^2 + \|\mathbf{b}(m)\|_2^2 - 2\langle \mathbf{a}(m), \mathbf{b}(m) \rangle$ and $|\langle \mathbf{a}(m), \mathbf{b}(m) \rangle| \leq \min\{S_x, S_y\}$, the lower and upper bounds on distance are given as:

$$S_x + S_y - 2 \min\{S_x, S_y\} \leq \|\mathbf{a}(m) - \mathbf{b}(m)\|_2^2 \leq S_x + S_y + 2 \min\{S_x, S_y\}. \quad (9)$$

This means that the upper bound on distances in the transform domain are restricted by sparsity levels.

In order to preserve distances, given the variance of the measurement noise σ_z^2 and sparsity level S_x ³, we have

³Sparsity level S_x obtained based on the desired distortion level imposed by sparse approximation.

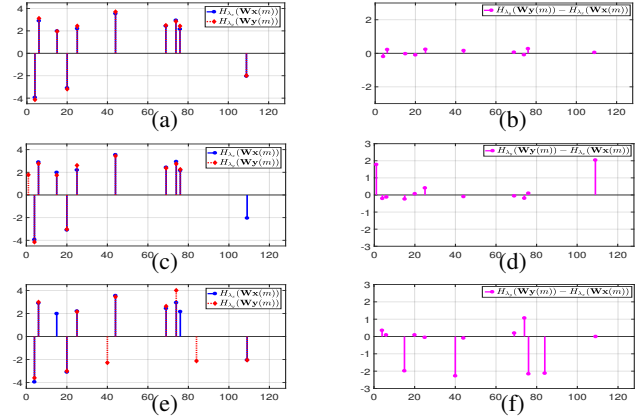


Fig. 3: Impact of the variance of measurement noise on the components of $H_{\lambda_x}(\mathbf{W}\mathbf{x}(m))$ and corresponding $H_{\lambda_y}(\mathbf{W}\mathbf{y}(m))$, for randomly selected $m \in [500]$, with $S_x = S_y = 10$, $N = L = 128$, $\sigma_x^2 = 1$. Fig. (3a), Fig. (3c) and Fig. (3e) represent the components of $H_{\lambda_x}(\mathbf{W}\mathbf{x}(m))$ and $H_{\lambda_y}(\mathbf{W}\mathbf{y}(m))$ with $\sigma_z^2 = 0.15$, $\sigma_z^2 = 0.2$ and $\sigma_z^2 = 0.4$, respectively. Fig. (3b), Fig. (3d) and Fig. (3f) indicate the corresponding errors.

to set $(1 - \epsilon) (N\sigma_z^2) \leq S_x + S_y - 2\langle \mathbf{a}(m), \mathbf{b}(m) \rangle \leq (1 + \epsilon) (N\sigma_z^2)$. So, the smallest integer sparsity level for approximation of noisy query $\mathbf{y}(m)$ is given as $S_y = \lfloor N\sigma_z^2 + 2\langle \mathbf{a}(m), \mathbf{b}(m) \rangle - S_x \rfloor$. Note that if $\mathbf{z} = \mathbf{0}$, then $S_x = S_y$ since in this condition $\mathbf{a}(m) = \mathbf{b}(m)$, therefore $\langle \mathbf{a}(m), \mathbf{b}(m) \rangle = S_x$. However, the above equation is implicit since $\mathbf{b}(m)$ depends on the sparsity level S_y , and vice versa. But, we can obtain it recursively. Fig. 2 depicts the required sparsity level S_y in order to preserve distances in the signal and transform domains. Due to the lack of space, we cannot provide our discussion about thresholding functions H_{λ_x} and H_{λ_y} as well as our analytic approach for distortion analysis. The more detailed analytic results will be presented in an extended version of this paper.

IV. AMBIGUIZATION ANALYSIS

In this section, we provide our results for ambiguitization of public storage. In order to make our ambiguitization approach more clear, first we discuss our model without ambiguitization noise. In Fig. 3, we illustrates impact of measurement noise on the sparse approximation with a hard-thresholding operator for small sparsity levels $S_x = S_y = 10$ ⁴. We see that small measurement noise just changes the magnitude of the components, while the indices of sparse components gradually miss by increasing the variance of the noise.

A. Ambiguitization in Sparsified Transform Domain

In this section, we propose a simple scheme for imposing ambiguitization at the server. In previous sections, we use \mathcal{T} as space of L -dimensional vectors in the transform domain. Let $\mathcal{U}, \mathcal{V} \subseteq \mathcal{T}$ be two subspaces such that $\mathcal{T} = \mathcal{U} + \mathcal{V}$. So, every vector $\mathbf{a} \in \mathcal{T}$ has at least one expression as $\mathbf{a} = \mathbf{u} + \mathbf{v}$, $\mathbf{u} \in \mathcal{U}, \mathbf{v} \in \mathcal{V}$. If we have $\mathcal{U} \cap \mathcal{V} = \{\mathbf{0}\}$, then every vector $\mathbf{a} \in \mathcal{T}$ has the *unique* expression $\mathbf{a} = \mathbf{u} + \mathbf{v}$, $\mathbf{u} \in \mathcal{U}, \mathbf{v} \in \mathcal{V}$ and we write $\mathcal{T} = \mathcal{U} \oplus \mathcal{V}$. Also, \mathcal{T} is called the direct sum of \mathcal{U} and \mathcal{V} . Now, let \mathcal{U} be the space of non-zero components

⁴The same sparsity levels are not optimal to preserve distances. We just intend to make sense of the influence of measurement noise on clean data.

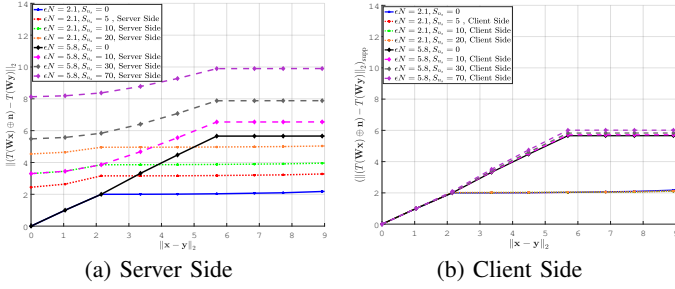


Fig. 4: Impact of ambiguization at (a) Server side and (b) Client side, for two different desired radius of $\|\mathbf{x} - \mathbf{y}\|_2 = \epsilon N = 2.1, 5.8$ in the original domain with different ambiguization noise sparsity levels S_{n_s} .

of \mathcal{T} and \mathcal{V} be the space of zero components of \mathcal{T} . The idea of our ambiguization method is to set ambiguization noise \mathbf{n} such that $\mathbf{n} \in \mathcal{V}$. Furthermore, since $(\mathcal{T}, \langle \cdot, \cdot \rangle)$ is an inner product space and $\mathcal{V} = \mathcal{U}^\perp \triangleq \{v \in \mathcal{V} : \langle u, v \rangle = 0, \forall u \in \mathcal{U}\}$, $\mathcal{T} = \mathcal{U} \oplus \mathcal{U}^\perp$ is orthogonal direct sum of \mathcal{U} and \mathcal{U}^\perp . It is clear that $\dim \mathcal{T} = L$, $\dim \mathcal{U} = S_x$ and $\dim \mathcal{U}^\perp = L - S_x$. We denote by S_{n_s} the sparsity level of ambiguization noise at the server. Note that $0 \leq S_{n_s} \leq L - S_x$.

Intuitively, when the number of choices increases, the server uncertainty increases, and so does the entropy measure proportionally to $\log_2 \binom{S_x + S_{n_s}}{S_x}$. By adding ambiguization noise \mathbf{n} to the zero components of the sparse representation of the clean data, we increase uncertainty at the server side. Clearly, we have a trade-off between utility and privacy. By setting sparsity levels such that $S_x \rightarrow 0$ and $S_{n_s} \rightarrow L$ we reach maximal privacy and loose our utility in the sense of a stable reconstruction, correct decision and also sparse representation. On the other hand, by choosing $S_x \rightarrow L$ and $S_{n_s} \rightarrow 0$ we reach maximal utility, however, we loose our privacy. However, it should be noted that authorized client shares a common secrecy with the data owner, i.e., he knows a set of secret positions in contrast to the server. Therefore, the addition of extra false components in the orthogonal positions does not harm the client as much as the server.

Fig. 4 depicts the effect of ambiguization on the distances at the server side and client side. Considering no ambiguization noise on the sparse representation of public data, the solid "blue" and "black" lines illustrate our desired preserved distances corresponding to $\|\mathbf{x} - \mathbf{y}\| = 2.1$ and $\|\mathbf{x} - \mathbf{y}\| = 5.8$, respectively. Hence, on the client side, we cannot recover any information about the NNs hash signals that are far apart, since our approach only preserves distances up to a radius, as determined by the sparsity level of the approximation. As it is shown, by imposing ambiguization noise, all distances from the server viewpoint go to constant value, i.e., all vectors in the public domain seem equally from the server standpoint. However, from the client side we can effectively preserve distances up to the designed radius, just by computing distances on non-zero components of the sparse representation of the query, i.e., on $\text{supp}(T_{\lambda_y}(\mathbf{W}\mathbf{y}(m)))$. Note that, regarding to our discussion in Sec. III, the upper bound on distances in the transform domain is restricted by sparsity levels.

In [2] the authors listed potential threats under the 'honest but curious' model. According to our previous discussions the threat S_1 in [2] precluded. In order to address threat




Fig. 5: Estimated PDFs of pairwise distances in transform domain with different sparsity ratios and privacy leak estimates $D(P_1 || P_2)$ for the corresponding cases.

S_2 . i.e., clustering the database vectors from $\mathbf{a}(m)$, $\forall m \in [M]$, we generate four 512-dimensional i.i.d. vectors with distribution $\mathcal{N}(\mathbf{0}, \mathbf{1})$ and 1000 512-dimensional i.i.d. vectors with distribution $\mathcal{N}(\mathbf{0}, \mathbf{0.1})$. Then we add each 250 (out of 1000) low variance vectors to the four high variance ones (schematically shown in Fig. 6). This results in a database of 1000 vectors adhering to four clusters. By defining sparsity ratios $\alpha_x = S_x/L$ and $\beta_x = S_{n_s}/L$, and denoting by P_{intra} and P_{inter} as probability density functions (PDFs) of 'intra-cluster' and 'inter-cluster' of distances, respectively, we intend to introduce a *Kullback-Leibler divergence privacy leak measure*. To this end, we define distribution P_1 as mixture of P_{intra} and P_{inter} , that is, $P_1 = \alpha_x P_{\text{intra}} + (1 - \alpha_x) P_{\text{inter}}$, $0 \leq \alpha_x \leq 1$. Also, we denote the Gaussian distribution $\mathcal{N}(\mu_2, \sigma_2^2)$ by P_2 , such that μ_2 and σ_2^2 are the mean and variance of P_1 , respectively. Therefore, the privacy leak measure can be defined by the Kullback-Leibler divergence (KLD) as $D(P_1 || P_2) = \alpha_x D(P_{\text{intra}} || P_2) + (1 - \alpha_x) D(P_{\text{inter}} || P_2)$, where $D(P_1 || P_2) = \mathbb{E}_{P_1} \left[\log \frac{P_1}{P_2} \right]$. Now, our privacy leak constraint can be expressed as $D(P_1 || P_2) \leq \delta$, where δ determines the allowable privacy leakage from database clustering. If all inter-class and intra-class distances follow the same distribution (close to Gaussian for large N or L), then no clustering algorithm can differentiate them on the server side. In the same way, the KLD $D(P_1 || P_2)$ will tend to zero.

In Fig. 5, we illustrate the estimated PDFs of pairwise distances in the transform domain. The 'solid lines' indicate P_{intra} and P_{inter} with different sparsity ratio β_x . As evident, by imposing ambiguization noise, the distributions P_{intra} and P_{inter} become unimodal Gaussian, therefore, the curious server cannot cluster database. Also, it is shown that by increasing the sparsity ratio of approximation, α_x , we have to increase the sparsity ratio of ambiguization, β_x , in order to keep our privacy protection. The 'dashed lines' indicate the corresponding Gaussian distribution fit to each 'solid' plot. These figures clarify our main intuition of defining the KL divergence as a privacy protection measure in the ambiguized transform domain. Fig. 6 visualizes how our model behaves in practice. Another interesting result is that k -means clustering will completely fail at the server, provided $D(P_1 || P_2) \leq \delta$. If the curious server runs the k -means scheme, the expected probability of correct clustering will be $\frac{1}{k}$, which shows maximal achievable Shannon entropy




Fig. 6: Schematic visualization of ambiguation on correlated clustered data used for Fig. 5.

for random guessing. In other words, the bin (label) assignment at the server is completely random and server cannot learn the structure of database.

Although in our framework we design the decoder to be in the private domain, in order to address threats S_3 and S_4 of [2] we generalize our framework and consider the *decoder* in the *public domain*. Now, we briefly discuss our solution for this scenario. In the public decoder scenario the data user (client side) just sends the positions of interest to the server and server returns lists of the vectors that have non-zero components for each position. We denote by $\mathcal{I} \subset [L]$ the set of indices that the user sends to the server such that $\text{card}(\mathcal{I}) = S_y + S_{n_q}$, where S_{n_q} is the number of ambiguated positions sent in addition to the S_y interested positions. For request position $l \in \mathcal{I}$, the server returns $\mathcal{L}^+(l) = \{m \in [M] : a_l(m) = +1\}$ and $\mathcal{L}^-(l) = \{m \in [M] : a_l(m) = -1\}$, where $\text{card}(\mathcal{L}^+(l)) \simeq \text{card}(\mathcal{L}^-(l)) \simeq 0.5M(\alpha_x + \beta_x)$. So, for a database consisting of $M = 10^9$ vectors with parameters $S_y + S_{n_q} = 16$ and $\alpha_x + \beta_x \approx 0.03125$, we need $M(\alpha_x + \beta_x)(S_y + S_{n_q}) \approx 62$ MB in order send back the lists to the user. To avoid the communication overhead, the server can communicate only the data for a short list after score aggregation. The data user refines the list by rejecting the results for ambiguated positions. To investigate the upper limit on clustering based on k-means for the example above, we investigate the ratio between the inter-class and intra-class distances for various parameters. Assuming $L = 256$, $S_x = 10$ and $\sigma_z^2 = 0.15$, we provide the average distances in TABLE I, where $\alpha_y = S_y/L$ and $\beta_y = S_{n_q}/L$. Defining \bar{d}_{diag} as the average of diagonal entries and \bar{d}_{off} as the average of off-diagonal entries for pairwise average distance matrices shown in TABLE I and setting $S_{n_s} = L - S_x$, we plot the ratio of $\bar{d}_{\text{off}} / \bar{d}_{\text{diag}}$ as a function of $(S_y + S_{n_q})/L = \alpha_y + \beta_y$ in Fig. 7. As evident, by increasing S_{n_q} from 0 to $L - S_y$, the ratio $\bar{d}_{\text{off}} / \bar{d}_{\text{diag}}$ converges to 1, i.e., we reach maximal privacy when the query discloses the positions to the server.

V. CONCLUSION

We have proposed a novel approach for privacy preserving identification based on a STC representation with ambiguation noise. Our scheme preserves distances up to a radius, as determined by the sparsity level of the approximation. However, for distances greater than our desired radius, the Euclidean distance between vectors in the transform domain

TABLE I: The upper limit of average distances between clusters when a user sends the positions to the server.

	$\beta_y \approx 0.04$			$\beta_y \approx 0.16$			$\beta_y \approx 1 - \alpha_y$					
\bar{d}_{diag}	0.9291	4.3568	4.1459	4.0621	1.8378	4.6563	4.4239	4.3569	4.4524	6.3278	5.9947	5.9989
\bar{d}_{off}	4.2247	0.9042	4.2149	4.2231	4.5523	1.8273	4.5442	4.5446	6.3276	4.4522	6.3211	6.3193
$\bar{d}_{\text{off}} / \bar{d}_{\text{diag}}$	4.1164	4.3912	0.92978	4.1577	4.3965	4.6747	1.8319	4.4322	5.9947	6.3211	4.4513	6.0044
$\bar{d}_{\text{off}} / \bar{d}_{\text{diag}}$	4.0754	4.4043	4.1618	0.9238	4.3832	4.7061	4.4547	1.8376	5.9989	6.3193	6.0044	4.4505
$\bar{d}_{\text{off}} / \bar{d}_{\text{diag}}$	5.4068	7.0727	6.7268	6.6932	9.4215	10.4913	10.2535	10.2365	22.0792	22.6286	22.5188	22.5135
$\bar{d}_{\text{off}} / \bar{d}_{\text{diag}}$	7.0680	5.4074	7.0168	6.9889	10.4819	9.4254	10.4600	10.4372	22.6286	22.0882	22.6180	22.5998
$\bar{d}_{\text{off}} / \bar{d}_{\text{diag}}$	6.7350	7.0452	5.4027	6.7650	10.2609	10.4687	9.4181	10.2751	22.5188	22.6180	22.0790	22.5328
$\bar{d}_{\text{off}} / \bar{d}_{\text{diag}}$	6.6914	7.0028	6.7587	5.4021	10.2311	10.4405	10.2758	9.4142	22.5135	22.5998	22.5328	22.0735




Fig. 7: Privacy protection of database when probe discloses the positions.

becomes independent of the distance between the vectors in the original domain. Our method is generic, i.e., we have imposed no restrictions on the input data. One of the main points illustrated by this study is that we can share our database in the public domain, while preserving data privacy and even ensure the reconstruction of data for authorized users in contrast to known methods. The results show that the curious server cannot cluster the samples in the database, provided the Kullback-Leibler privacy leak constraint is satisfied. Furthermore, we proposed a public decoding method by introducing a novel formulation of query mechanism based on sending non-zero positions along with the ambiguous positions to server.

ACKNOWLEDGMENT

B. Razeghi has been supported by the ERA-Net project ID_IoT No 20CH21_167534 and O. Taran by the SNF project No 200021_165672.

REFERENCES

- [1] S. Ferdowsi, S. Voloshynovskiy, D. Kostadinov, and T. Holotyak, "Fast content identification in high-dimensional feature spaces using sparse ternary codes," in *IEEE Int. Work. on Inf. Forensics and Security (WIFS)*, 2016, pp. 1–6.
- [2] B. Mathon, T. Furon, L. Amsaleg, and J. Bringer, "Secure and efficient approximate nearest neighbors search," in *1st ACM Work. on Inf. Hiding and Multi. Security*, 2013, pp. 175–180.
- [3] S. Rane and W. Sun, "An attribute-based framework for privacy preserving image querying," in *19th IEEE Int. Conf. on Image Processing (ICIP)*, 2012, pp. 2649–2652.
- [4] S. Ferdowsi, S. Voloshynovskiy, D. Kostadinov, and T. Holotyak, "Sparse ternary codes for similarity search have higher coding gain than dense binary codes," in *IEEE Int. Symp. on Inf. Theory (ISIT)*, 2017.
- [5] P. Boufounos and S. Rane, "Secure binary embeddings for privacy preserving nearest neighbors," in *IEEE Int. Work. on Inf. Forensics and Security (WIFS)*, 2011, pp. 1–6.
- [6] A. Iscen and T. Furon, "Group testing for identification with privacy," in *4th ACM Work. on Inf. Hiding and Multi. Security*, 2016, pp. 51–56.
- [7] H. Jégou, M. Douze, and C. Schmid, "On the burstiness of visual elements," in *IEEE Conf. on Comp. Vision and Pattern Recog. (CVPR)*, 2009, pp. 1169–1176.
- [8] F. Perronnin and C. Dance, "Fisher kernels on visual vocabularies for image categorization," in *IEEE Conf. on Comp. Vision and Pattern Recog. (CVPR)*, 2007, pp. 1–8.
- [9] H. Jégou, M. Douze, C. Schmid, and P. Pérez, "Aggregating local descriptors into a compact image representation," in *IEEE Conf. on Comp. Vision and Pattern Recog. (CVPR)*, 2010, pp. 3304–3311.
- [10] A. Babenko, A. Slesarev, A. Chigorin, and V. Lempitsky, "Neural codes for image retrieval," in *Europ. Conf. on Comp. Vision*, 2014, pp. 584–599.
- [11] S. Ravishanker and Y. Bresler, "Closed-form solutions within sparsifying transform learning," in *IEEE Int. Conf. on Acoustics, Speech and Signal Proc. (ICASSP)*, 2013, pp. 5378–5382.
- [12] —, "Learning sparsifying transforms," *IEEE Trans. on Signal Processing*, vol. 61, no. 5, pp. 1072–1086, 2013.
- [13] P. H. Schönemann, "A generalized solution of the orthogonal procrustes problem," *Psychometrika*, vol. 31, no. 1, pp. 1–10, 1966.
- [14] E. Kreyszig, *Introductory functional analysis with applications*. Wiley New York, 1989.