RECONSTRUCTION OF PRIVACY-SENSITIVE DATA FROM PROTECTED TEMPLATES

Shideh Rezaeifar, Behrooz Razeghi, Olga Taran, Taras Holotyak, Slava Voloshynovskiy

University of Geneva
{shideh.rezaeifar, behrooz.razeghi, olga.taran, taras.holotyak, svolos}@unige.ch

ABSTRACT

In this paper, we address the problem of data reconstruction from privacy-protected templates, based on recent concept of sparse ternary coding with ambiguization (STCA). The STCA is a generalization of randomization techniques which includes random projections, lossy quantization, and addition of ambiguization noise to satisfy the privacy-utility trade-off requirements. The theoretical privacy-preserving properties of STCA have been validated on synthetic data. However, the applicability of STCA to real data and potential threats linked to reconstruction based on recent deep reconstruction algorithms are still open problems. Our results demonstrate that STCA still achieves the claimed theoretical performance when facing deep reconstruction attacks for the synthetic i.i.d. data, while for real images special measures are required to guarantee proper protection of the templates.

Index Terms— Privacy, template protection, reconstruction, ambiguization, deep learning.

1. INTRODUCTION

Modern machine learning is based on the usage of massive data sets that often contain privacy-sensitive information. A similar problem exists with biometrics that are used in both private security systems granting access to various devices and services, and public security systems covering various surveillance and monitoring applications. In recent years, the advancement of personalized medicine applications also requires reliable privacy protection of genetic data and privacy sensitive clinical records. Despite the broad variety of these applications, machine learning tools are more often used to extract templates from the data that by itself does not guarantee their privacy protection against reconstruction attacks. At the same time, it is demonstrated that the original data can be reliably reconstructed from templates and non-linear representations extracted by both hand-crafted methods based on local descriptors [1] and deep representations [2]. Once successfully reconstructed, an adversary might use these data to impair both privacy and security in the above applications.

S. Rezaeifar and O. Taran are supported by the SNF project No 200021_182063, B. Razeghi is by the ERA-Net project ID_IoT No 20CH21_167534. S. Voloshynovskiy is a corresponding author.

The problem of template privacy protection in view of reconstruction attacks is well recognized and various generic methods were proposed such as fuzzy commitment schemes [3] and secure sketches [4], helper data based methods [5, 6], concealable template protection and robust hashing [7] as well as several practical methods in biometrics applications [2]. We do not pretend to be exhaustive in our overview and refer interesting readers to [8]. Recently, a concept of the STCA was proposed that combines and extends the encoding and randomization principles from the information-theoretic perspectives [9–11]. The STCA ensures the protection of both templates and queries in authentication and identification systems against the adversarial reconstruction and clustering [9, 10]. In contrast to binary templates, the STCA is based on sparse ternary encoding [12] of data ensuring the maximum information preservation for the authorized users, while minimizing the leakages for unauthorized users using a special ambiguization scheme based on an addition of noise to the zero components of ternary codes that do not contain any significant information. The authorized users benefit from the presence of degraded authentic data in a form of probe allowing the reliable reconstruction of the original features with a minimum loss of information. Thus, the STCA enables the verification and identification in the original space in contrast to the binarized templates, which face the loss of information, and yet enjoying fast search and optimal performance [9–11].

To the best of our knowledge, the STCA was not investigated under the deep adversarial reconstruction. Furthermore, the methods of template reconstruction based on the deep machine learning techniques considered in [1] and [2] were not investigated under the advanced privacy-preserving methods. Therefore, the goal of this paper is to make one step forward and consider the generalized reconstruction from the STCA protected data under adversarial deep reconstruction attacks. Along this way, we target to practically confirm the achievability of theoretical limits based on previously reported results [9, 10]. For this purpose, we will use synthetic data and analytically treatable feature extraction methods from one side and real data from another one. The adversarial attacks will be investigated in two settings of reconstruction from the protected template with all 'known model parameters' besides the randomization noise, that will be kept secret for the attacker, and 'unknown model parameters'.

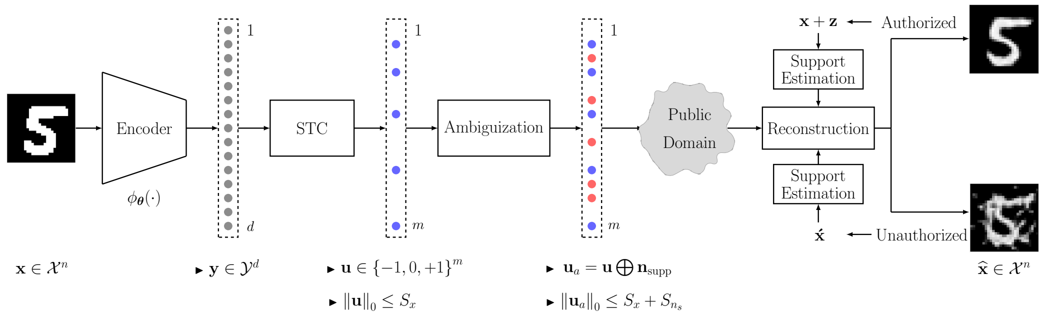

Fig. 1: General block diagram explaining the adversarial reconstruction of $\hat{\mathbf{x}}$ based on publicly available privacy protected template $\mathbf{u}_a$ based on the STCA applied to the template $\mathbf{y}$ extracted from $\mathbf{x}$ using an extractor $\phi_{\boldsymbol{\theta}}(\cdot)$.

Notations: X and $\mathbf{X}$ denote random variables and random vectors, while their realizations are denoted as x and $\mathbf{x}$, respectively.

Paper Organization: The problem formulation is given in Section 2. The reconstruction under the known model parameters is addressed in Section 3 and under the unknown parameters in Section 4. Section 5 presents the experimental results and Section 6 concludes the paper.

2. PROBLEM FORMULATION

Consider a signal $\mathbf{x} \in \mathcal{X}^n$, where $\mathcal{X} \subset \mathcal{R}$. A generic template extraction scheme is denoted as:

$$\mathbf{y} = \phi_{\boldsymbol{\theta}}(\mathbf{x}), \tag{1}$$

where $\phi_{\boldsymbol{\theta}}(\mathbf{x}) = \sigma_L(W_L ... \sigma_1(W_1\mathbf{x}))$ denotes any function in the form of a deep network parametrized by some learnable parameters $\boldsymbol{\theta} = \{W_1, \cdots, W_L\}$ with a set of element-wise non-linearities $\sigma_l(\cdot)$, $1 \leq l \leq L$, such as VGG_19 [13], AlexNet [14], etc. Furthermore, the template $\mathbf{y}$ can also be raw data, extracted features using any known hand crafted methods, aggregated local descriptors based on BoW, FV, VLAD [15–17]. Note that we do not impose any constraints on the sparsity of $\mathbf{x}$ in the direct domain or some transform domain.

In this work, we will focus on a simple model $\phi_{\boldsymbol{\theta}}(\mathbf{x}) = W\mathbf{x}$ to investigate the capacity of machine learning methods versus analytical closed-form solutions for reconstruction attacks against protected templates. Our setup is explained by the need to investigate the security of template protection rather than template extraction since the reconstruction from various features and templates is known to be successful according to [1, 2].

As privacy protection, we consider a generic privacy-preserving encoding based on random projections, quantization and addition of random noise that is integrated in the STCA framework [9, 10]:

$$\mathbf{u}_a = \varphi(A\mathbf{y}) \oplus \mathbf{n}_{\text{supp}}, \tag{2}$$

where A is a random $m \times d$ matrix, $\varphi(\cdot) : \mathcal{R} \to \{-1, 0, +1\}$ is a quantization operator representing an element-wise non-linearity and $\mathbf{n}_{\text{supp}}$ denotes the ambiguization noise that can be added to the orthogonal complement of space of $\varphi(A\mathbf{y})$. The sparse ternary representation $\mathbf{u} = \varphi(A\mathbf{y}) \in \{-1, 0, +1\}^m$ has S_x non-zero components, i.e., $\|\mathbf{u}\|_0 = S_x$, and protected template $\mathbf{u}_a$ has $S_x + S_{n_s}$ non-zero components, where $0 \leq S_{n_s} \leq m - S_x$ is the sparsity level of ambiguization noise for the public representations. Note that A may reduce, keep or extend the dimension of the vector $\mathbf{y}$, i.e., $m \gtreqless d$. The clean representation $\mathbf{u}$ will be used for adversarial training.

In this paper, our purpose is to evaluate practically the privacy-protection capacity of STCA based template protection by utilizing machine learning reconstruction tools. The general block diagram of our framework is depicted in Fig. 1. We investigate both synthetic and real data in our study. We consider two attacking scenarios assuming that: (a) $\varphi(\cdot)$, A and W are known to the attacker and $\mathbf{n}_{\text{supp}}$ is unknown and (b) $\varphi(\cdot)$, A, W, $\mathbf{n}_{\text{supp}}$ are unknown to the attacker. The goal of the attacker is to reconstruct $\hat{\mathbf{x}}$ as close as possible to $\mathbf{x}$ based on the protected template $\mathbf{u}_a$.

3. RECONSTRUCTION UNDER THE KNOWN MODEL PARAMETERS

In the most general case, the reconstruction under the known models (1) and (2) can be formulated as:

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \frac{1}{2} ||\mathbf{u}_a - \varphi(A\phi_{\boldsymbol{\theta}}(\mathbf{x}))||_2^2 + \lambda\Omega(\mathbf{x}), \tag{3}$$

where $||.||_2$ denotes a ℓ_2-norm, λ is a regularization parameter and $\lambda\Omega(\mathbf{x}) = -\log p(\mathbf{x})$. Unfortunately, it is not feasible to define a multidimensional pdf $p(\mathbf{x})$ in practice besides some rare exceptions of i.i.d. models and the models capturing sparsity.

To ensure privacy protection against adversarial reconstruction, we separate the impact of the imposed non-linearity by the network $\phi_{\boldsymbol{\theta}}(\cdot)$ on the extracted templates, from the STCA template protection capacity. Accordingly, we consider the network $\phi_{\boldsymbol{\theta}}(\mathbf{x}) = W\mathbf{x}$ with a constraint $W^TW = I$, that yields:

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \frac{1}{2} ||\mathbf{u}_a - \varphi(AW\mathbf{x})||_2^2 + \lambda\Omega(\mathbf{x}). \tag{4}$$

In the general case, assuming that $\varphi(\cdot)$ is differentiable with respect to $\mathbf{x}$, one can find the solution to (4) as:

$$\mathbf{x} \leftarrow \mathbf{x} - \eta \nabla_{\mathbf{x}} J(\mathbf{x}), \tag{5}$$

by denoting $J(\mathbf{x}) = \frac{1}{2}||\mathbf{u}_a - \varphi(AW\mathbf{x})||_2^2 + \lambda\Omega(\mathbf{x})$. However, the solution of (4) based on (5) faces several practical problems that we summarize below.

Problem 1 (the non-differentiability of $\varphi(\cdot)$): The STCA ternarization operator $\varphi(\cdot)$ is not differentiable with respect to $\mathbf{x}$. Therefore, one can envision several approaches to approximate $\varphi(\cdot)$ by some differentiable surrogate function, as for example by a hard-thresholding operator that preserves the same mutual information as for the ternarization operator for a range of sparsity levels S_x [18], or by considering a linear approximation, i.e., $\varphi(AW\mathbf{x}) = AW\mathbf{x}$, that yields to:

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \frac{1}{2}||\mathbf{u}_a - AW\mathbf{x}||_2^2 + \lambda\Omega(\mathbf{x}), \quad (6)$$

with $\Omega(\mathbf{x}) = ||\mathbf{x}||_2^2$. In this case, the solution reduces to:

$$\hat{\mathbf{x}} = ((AW)^T AW + \lambda I)^{-1}(AW)^T \mathbf{u}_a. \quad (7)$$

Problem 2 (model prior for real data): The above assumed ℓ_2-norm regularizer works only for the synthetic i.i.d Gaussian data. In the case of real images, it is too restrictive. The class of sparsification priors is also relatively restrictive in view of a single overcomplete shallow representation. Instead, recent works [19,20] suggested using a generative model $\mathbf{x} = g_{\boldsymbol{\theta}_G}(\mathbf{z})$, where $g_{\boldsymbol{\theta}_G}(\cdot)$ is a generator of GAN or decoder of VAE trained on corresponding data $\{\mathbf{x}_i\}_{i=1}^N$, where N denotes the number of training samples. In this case, the reconstruction problem reduces to:

$$\hat{\mathbf{z}} = \arg\min_{\mathbf{z}} \frac{1}{2}||\mathbf{u}_a - \varphi(AWg_{\boldsymbol{\theta}_G}(\mathbf{z}))||_2^2 + \lambda\Omega(\mathbf{z}), \quad (8)$$

and $\hat{\mathbf{x}} = g_{\boldsymbol{\theta}_G}(\hat{\mathbf{z}})$ under the differentiability or surrogate replacement of non-differentiable $\varphi(\cdot)$.

4. RECONSTRUCTION UNDER THE UNKNOWN MODEL PARAMETERS

The adversarial reconstruction under the model with unknown parameters is of more practical interest. Instead of assuming the complete knowledge of the model and its differentiability, the adversary has access to the training data $\{\mathbf{x}_j, \mathbf{u}_j\}_{j=1}^M$ or can use the model $\mathbf{u} = \varphi(A\mathbf{y})$ as a *black box* for the given inputs $\{\mathbf{x}_j\}_{j=1}^M$ to compute the templates $\{\mathbf{u}_j\}_{j=1}^M$.

The recovery problem reduces to the training of a reconstruction deep network $f_{\boldsymbol{\psi}}(\mathbf{u})$:

$$\boldsymbol{\psi}^* = \arg\min_{\boldsymbol{\psi}} \frac{1}{M}\sum_{j=1}^{M} \mathcal{L}(f_{\boldsymbol{\psi}}(\mathbf{u}_j), \mathbf{x}_j). \quad (9)$$

The trained reconstruction network represents a deep decoder that is applied directly to a privacy-protected template $\mathbf{u}_a$ by producing the recovered image $\hat{\mathbf{x}} = f_{\boldsymbol{\psi}^*}(\mathbf{u}_a)$. One can also envision a setup of training on protected templates $\mathbf{u}_a$ with different ambiguization levels of $\mathbf{u}$ that is out of the scope of this paper.

5. EXPERIMENTAL RESULTS

The privacy protection power of STCA originates from the lossy ternary quantization induced by $\varphi(\cdot)$ and the addition of the ambiguization noise $\mathbf{n}_{\text{supp}}$ to the support complement of sparse ternary approximation. Clearly, the imposed lossy ternarization has an impact on both the authorized and unauthorized users that prevents the perfect recovery of $\mathbf{x}$. We can measure how much information is lost and how much is preserved, in the terms of distortion level and encoding rate. Moreover, we can link it to the classical Shannon rate-distortion theory. The intelligently designed ambiguization scheme proposed by the STCA ensures accurate reconstruction for the authorized users while prohibits an accurate reconstruction for the unauthorized users. This originates from the fact that the authorized users can estimate the correct support of the data using their noisy data $\mathbf{x} + \mathbf{z}$, while the unauthorized users have no knowledge to "unlock" the protected template.

Consequently, we will investigate: (a) the link to the rate-distortion function, (b) the reconstruction based on the pseudo-inverse (7) and (c) the reconstruction based on the trained network (9). To evaluate the achievability of theoretical limits, we will first validate our results on synthetic data and then extend them to real images. As the synthetic data, we used i.i.d. Gaussian samples $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \sigma_X^2 I_n)$ and mappers $A_{i,j} \sim \mathcal{N}(0, 1/\sqrt{n})$ and DCT transform as W. As the real images, we have used MNIST data set. Although, this dataset is relatively simple however it is very useful to assess the quality of reconstruction in terms of unique recognition of digits by human thus avoiding any subjective factors of quality evaluation. The mappers A and W are the same as for the synthetic data.

We present three series of tests: (i) the theoretical limits of reconstruction from the STC representations in terms of achieving the Shannon lower bound on rate-distortion [21] is investigated in Figures 2a and 2b, (ii) the capability of adversary to reconstruct from the protected templates $\mathbf{u}_a$ with different levels of sparsity is depicted in Figures 2c and 2d and (iii) the accuracy of reconstruction from the protected templates for authorized users with access to the noisy data is shown in Figures 2e and 2f.

All results are shown for the reconstruction based on the pseudo-inverse (7) with $\lambda = 0$ and two types of deep reconstruction networks. The reconstruction network (9) for the synthetic data consists of: Linear(529)$\rightarrow$Tanh $\rightarrow$ Reshape2D $\rightarrow$ Conv2D (channels=32, kernel=7) $\rightarrow$ ReLu $\rightarrow$ Conv2D (channels=16, kernel=5)$\rightarrow$Tanh$\rightarrow$Conv2D (channels=8, kernel=3) $\rightarrow$ ReLu $\rightarrow$ Conv2D (channels=1, kernel=3) $\rightarrow$ Tanh. The reconstruction network for the real data consists of: Linear(784) $\rightarrow$ ReLu $\rightarrow$ Reshape2D $\rightarrow$ Conv2D (channels=32, kernel=5) $\rightarrow$ ReLu $\rightarrow$ Conv2D(channels=16, kernel=5) $\rightarrow$ ReLu$\rightarrow$Conv2D (channels=1, kernel=5)$\rightarrow$ReLu.

The obtained results demonstrate that the pseudo-inverse reconstruction and deep net reconstruction for the i.i.d. synthetic data have very similar performance while for the real data the deep reconstruction benefits from the presence of structured training data and produces perceptually more pleasant results as shown in Figure 3. The pseudo-inverse reconstruction does not use data priors and the produced

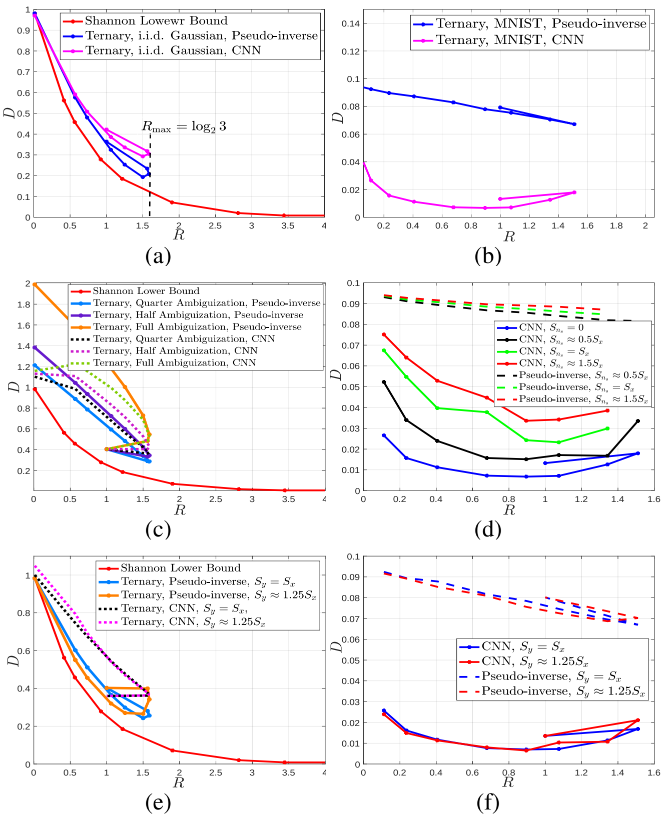

Fig. 2: Comparison of Rate-Distortion behavior on the synthetic and real data sets. (a), (c), (e): comparison of pseudo-inverse and CNN reconstructions on synthetic data set. (b), (d), (f): comparison of pseudo-inverse and deep reconstructions on MNIST data set. (a), (b): $R(D)$ curve limits. (c), (d): $R(D)$ curve for unauthorized users with different ambiguization levels. (e), (f): $R(D)$ curve for authorized users with different sparsity levels S_y and considering noisy measurement $\mathbf{x} + \mathbf{z}$, such that $\sigma_Z^2 = 0.25\,\sigma_X^2$.

results are quite noisy.

It is interesting to note that both types of reconstruction strategies approach the Shannon limit on the synthetic data for the rates less than $\log_2 3 = 1.585$, which is the maximum achievable rate by the ternary encoding used in the STCA. This behaviour is depicted in Fig. 2a. However, the pseudo-inverse reconstruction produces the results closer to the Shannon lower bound. This can be explained that the encoding matrix is well defined and known while in the case of CNN reconstruction the inverse mapping is learned from the training i.i.d data only.

The ambiguization has a strong influence on the adversary reconstruction ability according to Fig. 2c and Fig. 2d that is clearly reflected by the increase of reconstruction distortion with the increase of the ambiguiszation S_{n_s}. However, in the case of real data, the structure in the data makes it possible for the unauthorized user to reconstruct from the protected template as shown in Fig. 2d. Several examples of the comparison between the pseudo-inverse reconstruction and those

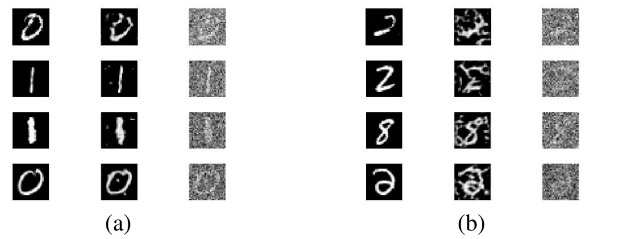

Fig. 3: Comparison of deep and pseudo-inverse reconstructions on MNIST images. (a): $S_x = 50$, $S_{n_s} = 25$, (b): $S_x = 50$, $S_{n_s} = 100$. First column of each panel corresponds to the original data, second and third ones correspond to reconstructed data based on deep network and pseudo-inverse, respectively.

based on the trained network are shown in Fig. 3. As expected, increasing the ambiguization makes the recognition more difficult and sometimes ambiguous (Fig. 3b). In practice, one should find the correct ambiguization factor to a particular type of data. Finally, the authorized user can unlock the ambiguization and approach the theoretical limit in contrast to the adversary (Fig. 2c, 2d) as shown in Figures 2e and 2f. This confirms the main research hypothesis in the paper.

6. CONCLUSIONS

We studied the capability of STCA for privacy-protection data release. The validates shown that for synthetic data, the STCA provides the theoretically achievable limits under both known and unknown model based adversarial reconstruction. Moreover, neglection of the ternarization non-linearity and usage of a simple pseudo-inverse does not lead to any drop in performance in comparison to deep non-linear reconstruction. An interesting point is that the deep reconstruction algorithm does not benefit from i.i.d data in comparison to simple projection matrix knowledge used in the pseudo-inverse.

The analysis on the real images established that the difference between the accuracy of reconstruction produced by the linear pseudo-inverse and deep reconstruction is significant besides the fact that the pseudo-inverse uses priors about the models of template extraction and template protection and deep reconstruction is based solely on the training data. We consider two major factors that impact this difference: (i) the fact that the linear pseudo-inverse does not use any prior about the statistics of images while the deep reconstruction can learn data manifold implicitly from the training data, (ii) the non-linear deep reconstruction also overcomes the problem of differentiability of the ternarization operator while the linear pseudo-inverse used in this paper uses the linear approximation. We believe that these factors should be covered in future research. In particular, the lack of reliably image priors can be addressed by the generative model in the formulation (8) under the linear ternarization approximation and compared with the solution based on (9). An additional research problem to be addressed is an *adversarial training* (9), when the various ambiguizations $\mathbf{u}_a$ should replace the clean version $\mathbf{u}$ similarly to the denoising auto-encoder training [22].

7. REFERENCES

[1] Alexey Dosovitskiy and Thomas Brox, "Inverting visual representations with convolutional networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 4829–4837.

[2] Guangcan Mai, Kai Cao, C YUEN Pong, and Anil K Jain, "On the reconstruction of face images from deep face templates," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.

[3] Yevgeniy Dodis, Leonid Reyzin, and Adam Smith, "Fuzzy extractors: How to generate strong keys from biometrics and other noisy data," in *International conference on the theory and applications of cryptographic techniques*. Springer, 2004, pp. 523–540.

[4] Julien Bringer, Hervé Chabanne, Gerard Cohen, Bruno Kindarji, and Gilles Zemor, "Theoretical and practical boundaries of binary secure sketches," *IEEE Transactions on Information Forensics and Security*, vol. 3, no. 4, pp. 673–683, 2008.

[5] Evgeny Verbitskiy, Pim Tuyls, Dee Denteneer, and Jean-Paul Linnartz, "Reliable biometric authentication with privacy protection," in *Proc. 24th Benelux Symposium on Information theory*, 2003, p. 19.

[6] Tanya Ignatenko and Frans MJ Willems, "Biometric systems: Privacy and secrecy aspects," *IEEE Transactions on Information Forensics and security*, vol. 4, no. 4, pp. 956, 2009.

[7] Yagiz Sutcu, Husrev Taha Sencar, and Nasir Memon, "A secure biometric authentication scheme based on robust hashing," in *Proceedings of the 7th workshop on Multimedia and security*. ACM, 2005, pp. 111–116.

[8] Karthik Nandakumar and Anil K Jain, "Biometric template protection: Bridging the performance gap between theory and practice," *IEEE Signal Processing Magazine*, vol. 32, no. 5, pp. 88–100, 2015.

[9] Behrooz Razeghi, Slava Voloshynovskiy, Dimche Kostadinov, and Olga Taran, "Privacy preserving identification using sparse approximation with ambiguization," in *IEEE International Workshop on Information Forensics and Security (WIFS)*, Rennes, France, December 2017, pp. 1–6.

[10] Behrooz Razeghi and Slava Voloshynovskiy, "Privacy-preserving outsourced media search using secure sparse ternary codes," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Calgary, Alberta, Canada, April 2018.

[11] Behrooz Razeghi, Slava Voloshynovskiy, Sohrab Ferdowsi, and Dimche Kostadinov, "Privacy-preserving identification via layered sparse code design: Distributed servers and multiple access authorization," in *26th European Signal Processing Conference (EUSIPCO)*, Rome, Italy, September 2018.

[12] Sohrab Ferdowsi, Sviatoslav Voloshynovskiy, Dimche Kostadinov, and Taras Holotyak, "Sparse ternary codes for similarity search have higher coding gain than dense binary codes," in *IEEE International Symposium on Information Theory (ISIT)*, Aachen, Germany, 2017.

[13] Karen Simonyan and Andrew Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.

[14] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.

[15] Hervé Jégou, Matthijs Douze, and Cordelia Schmid, "On the burstiness of visual elements," in *IEEE Conf. on Comp. Vision and Pattern Recog. (CVPR)*, 2009, pp. 1169–1176.

[16] Florent Perronnin and Christopher Dance, "Fisher kernels on visual vocabularies for image categorization," in *IEEE Conf. on Comp. Vision and Pattern Recog. (CVPR)*, 2007, pp. 1–8.

[17] Hervé Jégou, Matthijs Douze, Cordelia Schmid, and Patrick Pérez, "Aggregating local descriptors into a compact image representation," in *IEEE Conf. on Comp. Vision and Pattern Recog. (CVPR)*, 2010.

[18] Sohrab Ferdowsi, "Learning to compress and search visual data in large-scale systems," 2019.

[19] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G Dimakis, "Compressed sensing using generative models," *arXiv preprint arXiv:1703.03208*, 2017.

[20] Piotr Bojanowski, Armand Joulin, David Lopez-Paz, and Arthur Szlam, "Optimizing the latent space of generative networks," *arXiv preprint arXiv:1707.05776*, 2017.

[21] Thomas M Cover and Joy A Thomas, *Elements of information theory*, John Wiley & Sons, 2012.

[22] Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, and Pierre-Antoine Manzagol, "Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion," *Journal of Machine Learning Research*, vol. 11, pp. 3371–3408, 2010.