

# A GAN-Based Image Transformation Scheme for Privacy-Preserving Deep Neural Networks

Warit Sirichotedumrong and Hitoshi Kiya

*Department of Computer Science, Tokyo Metropolitan University, Hino-shi, Tokyo, Japan*

**Abstract**—We propose a novel image transformation scheme using generative adversarial networks (GANs) for privacy-preserving deep neural networks (DNNs). The proposed scheme enables us not only to apply images without visual information to DNNs, but also to enhance robustness against ciphertext-only attacks (COAs) including DNN-based attacks. In this paper, the proposed transformation scheme is demonstrated to be able to protect visual information on plain images, and the visually-protected images are directly applied to DNNs for privacy-preserving image classification. Since the proposed scheme utilizes GANs, there is no need to manage encryption keys. In an image classification experiment, we evaluate the effectiveness of the proposed scheme in terms of classification accuracy and robustness against COAs.

**Index Terms**—Deep neural network, generative adversarial network, privacy-preserving, visual protection

## I. INTRODUCTION

Deep neural networks (DNNs) have become one of the most powerful approaches for solving complex problems for many applications [1]–[3], such as for computer vision, biomedical systems, natural language processing, and speech recognition. DNNs gain an advantage by using a large amount of data to extract representations of relevant features, so the performance is significantly improved [4]. However, DNNs have been deployed in security-critical applications, such as facial recognition, biometric authentication, and medical image analysis.

Instead of using local servers, it is very popular for data owners to use cloud servers to compute a huge amount of data, but several data security issues have to be considered for utilizing DNNs, such as data privacy, data leakage, and unauthorized data access. Moreover, DNNs can be vulnerable if adversaries carry out model inversion attacks [5], [6] to obtain trained data from the trained model. Hence, privacy-preserving DNNs have become an urgent challenge. In this paper, we focus on protecting visual information of images before uploading the images to cloud environments.

Various perceptual image encryption methods have been proposed to protect visual information of images [7]–[15]. In contrast to information theory-based encryption (like RSA and AES), images encrypted by the perceptual encryption methods can be directly applied to some image processing algorithms.

A few previous works [10], [14], [15] have been proposed not only to protect the visual information of images but also to be applicable to DNNs. However, the visual information of

This work was partially supported by Grant-in-Aid for Scientific Research(B), No.17H03267, from the Japan Society for the Promotion Science.

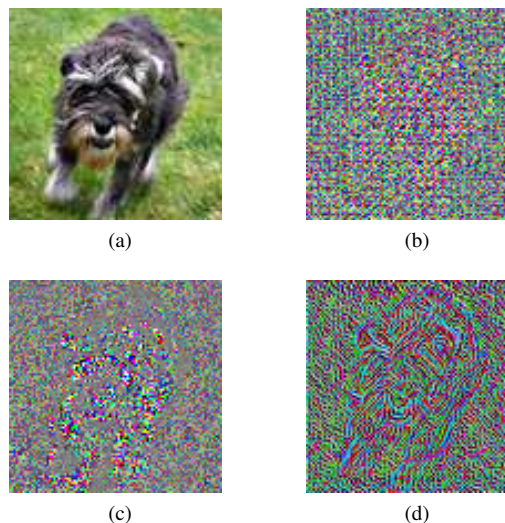


Fig. 1: Example of encrypted images. (a) Original image ( $U \times V = 96 \times 96$ ). (b) Tanaka's scheme [10] (c) Pixel-based encryption [14], [15]. (d) Proposed scheme.

images encrypted by using these methods have been pointed out to be reconstructed if adversaries carry out DNN-based attacks [15]–[17].

This paper proposes a novel image transformation scheme for privacy-preserving DNNs that enables us not only to apply images without visual information to DNNs, but also to enhance robustness against COAs. Since the proposed scheme utilizes the framework of generative adversarial networks (GANs), there is no need to manage encryption keys. In an image classification experiment, the proposed scheme is demonstrated to enable us not only to maintain high classification accuracy as that of conventional methods but also to have stronger robustness against COAs including DNN-based attacks than conventional ones.

## II. RELATED WORK

### A. Image Encryption for Privacy-Preserving DNNs

Various perceptual image encryption methods have been proposed to protect visual information of plain images [7]–[15]. Visually-protected images generated by using an encryption method consist of pixels, as shown in Figs. 1(b) and 1(c). Therefore, the encrypted images can be directly applied to image processing algorithms. Some of encryption methods

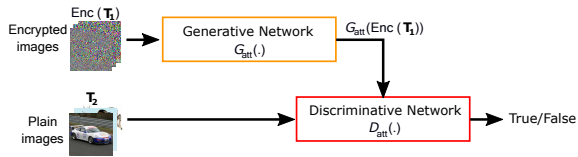


Fig. 2: The framework of GAN-based attack (GA)

have been studied for applying encrypted images to traditional machine learning algorithms, such as support vector machine (SVM), under the use of the kernel trick [18], [19], but they cannot be applied to DNNs [15].

There are three encryption methods [10], [14], [15] for privacy-preserving DNNs. The first method is Tanaka's scheme [10], which utilizes an adaptation network prior to DNNs to reduce the influence of image encryption. The second [14] is a pixel-based method, which enables data augmentation in the encrypted domain, under the use of a common encryption key to every image. The third one [15] enables us to assign an independent encryption key to every image; therefore, there is no need to manage keys. However, these methods do not consider robustness against DNN-based attacks [16]. Hence, this paper aims to produce visually-protected images having robustness against various COAs including DNN-based ones, while maintaining high classification accuracy.

### B. Robustness against Ciphertext-only Attacks

In this paper, we focus on robustness against COAs, such as brute-force attacks and DNN-based attacks. Conventional encryption methods [10], [14], [15] have large key spaces, so they are robust against brute-force attacks. However, they are pointed out not to have enough robustness against DNN-based attacks [15].

There are two types of DNN-based attacks: unshared weight attack (UWA) [16] and GAN-based attack (GA) [17]. Although the pixel-based encryption with independent encryption keys [15] has stronger robustness against UWA than that with a common key, robustness against GA has not been evaluated yet. The detail of GA is described in this section.

Figure 2 illustrates the framework of GA which consists of one generative network and one discriminative network. The generative network ( $G_{att}(\cdot)$ ) reconstructs visual information on plain images from the encrypted images while the discriminative network ( $D_{att}(\cdot)$ ) distinguishes the difference between reconstructed images and the plain ones. The network architectures of  $D_{att}(\cdot)$  and  $G_{att}(\cdot)$  used in this paper are illustrated in Fig. 3 and Fig. 4, respectively.

To train GA, a set of training images  $\mathbf{T}$  is equally divided into two image sets:  $\mathbf{T}_1$ , and  $\mathbf{T}_2$ , and then  $\mathbf{T}_1$  is encrypted by using an encryption scheme ( $Enc(\cdot)$ ). Eventually,  $Enc(\mathbf{T}_1)$  and  $\mathbf{T}_2$  are employed for training  $G_{att}(\cdot)$  and  $D_{att}(\cdot)$  (See Fig. 2).

## III. PROPOSED SCHEME

### A. Overview

Figure 5(a) illustrates the use of a transformation network ( $h_p(\cdot)$ ) for protecting training and testing images in DNNs. In




Fig. 3: Discriminative network  $D_{att}(\cdot)$

order to obtain  $h_p(\cdot)$ , as depicted in Fig. 5(b), the proposed scheme has two processes: preliminary visual protection and transformation network training. In the preliminary visual protection process, a set of preliminary protected images ( $\mathbf{P}$ ) is generated from a pre-trained model ( $h_\theta(\cdot)$ ) by using a set of training images ( $\mathbf{X}$ ) and the corresponding labels ( $\mathbf{Y}$ ).  $\mathbf{X}$  with  $\mathbf{Y}$ , and  $\mathbf{P}$  are utilized for training  $h_p(\cdot)$  which is used for generating a set of visually-protected images ( $\mathbf{X}_p$ ). Namely,  $\mathbf{X}_p = h_p(\mathbf{X})$ .

Hence, images used for training and testing DNNs are protected by using  $h_p(\cdot)$ , as shown in Fig. 5(a).

### B. Preliminary Visual Protection Process

We aim to generate a protected image ( $x_p$ ) from a plain image ( $x$ ) with a label ( $y$ ) by adding some perceptible noise  $\Delta$ ; therefore,  $x_p = x + \Delta$ . Let  $x_i$  and  $y_i$  denote  $i$ -th image with a label, where  $x_i \in \mathbf{X}$  and  $y_i \in \mathbf{Y}$ . In contrast to adversarial examples [20], protected images are generated by minimizing the loss with respect to input  $x$  with  $\Delta$  as given by

$$\underset{\Delta \in \mathcal{P}}{\text{minimize}} \frac{1}{m} \sum_{i=1}^m \mathcal{L}(h_\theta(x_i + \Delta), y_i), \quad (1)$$

where  $\mathcal{P}$  is a perturbation set, and  $\mathcal{L}(\cdot)$  is a loss function. The perturbation set  $\mathcal{P}$ , which is used for protecting visual information on plain images, can be written as

$$\mathcal{P} = \{\Delta : \|\Delta\|_\infty \leq \epsilon\}, \quad (2)$$

where  $\epsilon$  is the magnitude of the perturbation. The proposed method is inspired by adversarial attacks called projected gradient descent (PGD). We minimize the loss, but the adversarial attack maximizes it. At the  $(t+1)^{th}$  iteration, the protected image ( $x_p^{t+1}$ ) is

$$x_p^{t+1} = \prod_{x+\Delta} (x_p^t + \alpha \text{sign}(\nabla_x (-\mathcal{L}(\theta, x, y))), \quad (3)$$

where  $\alpha$  is a step size. In this paper,  $\Delta$  is iteratively added to an image for 50 iterations.

Therefore,  $h_\theta(\cdot)$  is expected to precisely classify the protected images, although the images have no visual information.

### C. Transformation Network

To obtain a transformation network  $h_p(\cdot)$ , we utilize unpaired image-to-image translation using cycle-consistent adversarial networks (CycleGAN) [21] to convert an image from the plain domain to the visually-protected domain.




Fig. 4: Architectures of generative network  $G_{att}(\cdot)$ .

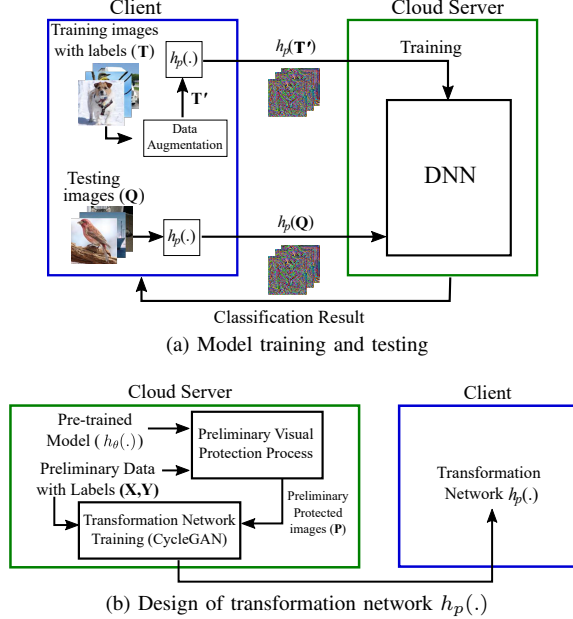


Fig. 5: Proposed scheme for privacy-preserving DNNs.




Fig. 6: CycleGAN architecture

1) *Formulation*: Let  $\mathbb{A}$  and  $\mathbb{B}$  denote the plain domain and the visually-protected domain, respectively. As shown in Fig. 6, CycleGAN consists of two generative networks:  $G_{AB} : \mathbb{A} \rightarrow \mathbb{B}$ , and  $G_{BA} : \mathbb{B} \rightarrow \mathbb{A}$ , which transform an image between two domains. In addition, there are two discriminative networks  $F_A$  and  $F_B$ .  $F_A$  discriminates the difference between  $x$  and  $G_{BA}(x_p)$  while  $F_B$  distinguishes the difference between  $x_p$  and  $G_{AB}(x)$ . The objective of training CycleGAN is to obtain  $G_{AB}$  which transforms plain images to visually-protected images, namely,  $h_p(\cdot) = G_{AB}$ .

2) *Loss Function*: According to [21], the CycleGAN utilizes two loss functions: adversarial loss ( $\mathcal{L}_{ad}$ ), and cycle consistency loss ( $\mathcal{L}_{cyc}$ ). Therefore, the full objective function

( $\mathcal{L}_{gan}$ ) can be expressed as

$$\begin{aligned} \mathcal{L}_{gan}(G_{AB}, G_{BA}, F_A, F_B) = & \lambda \mathcal{L}_{ad}(G_{BA}, F_A, x, x_p) \\ & + \lambda \mathcal{L}_{ad}(G_{AB}, F_B, x_p, x) \\ & + \mathcal{L}_{cyc}(G_{BA}, G_{AB}), \end{aligned} \quad (4)$$

where  $\lambda$  is utilized for weighting the relation between  $\mathcal{L}_{ad}$  and  $\mathcal{L}_{cyc}$ .

In order to ensure that images generated by  $h_p(\cdot)$  have no visual information while maintaining high classification accuracy, we propose a new cycle consistency loss which can be written as

$$\mathcal{L}_{cyc} = \gamma_1 \cdot \mathcal{L}_p + \gamma_2 \cdot \mathcal{L}_c + \gamma_3 \cdot \mathcal{L}_r, \quad (5)$$

where  $\mathcal{L}_p$ ,  $\mathcal{L}_c$ , and  $\mathcal{L}_r$  denote a perceptual loss, classification loss, and pixel-wise reconstruction loss, respectively.

We first aim to maximize the perceptual loss ( $\mathcal{L}_p$ ) [22] to protect the visual information. Since the CycleGAN consists of two cycles:  $\mathbb{A} \rightarrow \mathbb{B}$ , and  $\mathbb{B} \rightarrow \mathbb{A}$ . The perceptual loss of our CycleGAN,  $\mathcal{L}_p$ , can be expressed as

$$\mathcal{L}_p = \mathcal{L}_{feat}(x, G_{AB}(x)) + \mathcal{L}_{feat}(x, G_{BA}(x_p)). \quad (6)$$

In accordance with [22], the feature loss ( $\mathcal{L}_{feat}$ ) is calculated by

$$\mathcal{L}_{feat}(x_i, \hat{x}_i) = \frac{1}{C_k H_k W_k} \|\phi_k(\hat{x}_i) - \phi_k(x_i)\|_2^2, \quad (7)$$

where  $\phi_k(x_i)$  is a feature map with a size of  $C_k \times H_k \times W_k$ , obtained by using the  $k^{th}$  layer of a network when image  $x_i$  is fed [22]. Note that  $\hat{x}_i$  is a transformed image generated by  $G_{AB}(\cdot)$  and  $G_{BA}(\cdot)$ . In this paper, we use VGG16 [23], which is pre-trained with ImageNet, as a network for  $\phi_k(x_i)$ . Note that we utilize the feature maps of the second ReLU fuction of VGG16.

Moreover, to maintain high classification accuracy, we solve the minimization problem of the classification loss ( $\mathcal{L}_{cls}$ ) by utilizing  $h_\theta$  from Section III-B, where  $\mathcal{L}_{cls}$  is given by the cross entropy loss. The classification loss of the CycleGAN ( $\mathcal{L}_c$ ) is

$$\mathcal{L}_c = \mathcal{L}_{cls}(h_\theta(G_{AB}(x), y)) + \mathcal{L}_{cls}(h_\theta(G_{BA}(x_p), y)) \quad (8)$$

As in [21], we also utilize the following reconstruction loss ( $\mathcal{L}_r$ ).

$$\mathcal{L}_r = \|G_{BA}(G_{AB}(x)) - x\|_2^2 + \|G_{AB}(G_{BA}(x_p)) - x_p\|_2^2. \quad (9)$$

In an experiment, we employ the architecture of U-Net [24] for  $G_{AB}$  and  $G_{BA}$ , and the architectures of  $F_A$  and  $F_B$  are the same as in [21]

TABLE I: Image classification accuracy and average structural similarity (SSIM) values of reconstructed images by DNN-based attacks. (N/A: not available)

Encryption Scheme	SSIM values		Accuracy (%)	
	UWA [16]	GA [17]	CIFAR -10	CIFAR -100
Proposed Scheme	N/A	<b>0.0956</b>	90.73	67.36
Pixel-based [14], [15]	Same	0.1715	91.29	67.17
	Different	0.0425	91.14	67.28
Tanaka's Scheme [10]	N/A	0.7152	88.96	68.13
Plain images	N/A	N/A	95.05	77.24

#### IV. EXPERIMENTS

Image classification experiments were carried out to confirm the effectiveness of the propose scheme.

##### A. Dataset

We employed two image datasets: CIFAR-10 and CIFAR-100, for training and evaluating the proposed scheme. The first one is CIFAR-10 dataset which contains  $32 \times 32$ -pixel color images and consists of 50K training images [25]. Namely, the transformation network  $h_p(\cdot)$  was trained with CIFAR-10 dataset.

To evaluate the classification performance, we employed CIFAR-10 and CIFAR-100 [25], which contain  $32 \times 32$ -pixels color images and consists of 50K training images and 10K test images in 10 classes and 100 classes, respectively. We utilized CIFAR-100, to prove that the proposed scheme is applicable to any datasets, even when the model was trained by using CIFAR-10.

For training and testing two DNN-based attacks: UWA [16] and GA [17], we employed STL-10 dataset [26], which contains  $96 \times 96$ -pixel color images and consists of 5K training images and 8K testing images.

##### B. Training Transformation Network

VGG13 [23] with batch normalization was trained by using CIFAR-10 so that this model was considered as  $h_\theta(\cdot)$ . Then, we obtained preliminary protected images (P) for training the CycleGAN with the proposed loss function in (5).

CycleGAN was trained for 5000 epochs by using the Adam optimizer [27] with a learning rate of 0.0002, and a momentum parameters of  $\beta_1 = 0.5$  and  $\beta_2 = 0.999$ , where  $\lambda = 0.4$ ,  $\gamma_1 = -1$ ,  $\gamma_2 = 0.4$ , and  $\gamma_3 = 0.9$ . As a result,  $h_p(\cdot)$  was obtained and used for generating visually-protected images. An example of images generated by  $h_p(\cdot)$  is illustrated in Fig. 1(d), where Fig. 1(a) is the original one.

##### C. Robustness against Attacks

1) *Training Conditions*: GA was trained for 100 epochs by using the Adam optimizer [27] with a learning rate of 0.0002, a momentum parameter of  $\beta = 0.5$ , and a batch size of 64. For reconstructing images encrypted by Tanaka's scheme and the pixel-based method [14], [15], we employed the network architectures in Figs. 4(a) and 4(b), respectively. U-Net [24]

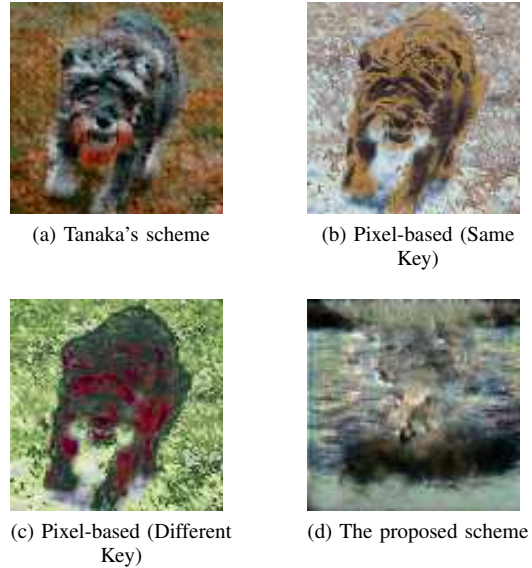


Fig. 7: Example of reconstructed images.

was utilized as  $G_{att}(\cdot)$  in order to reconstruct transformed images generated by the proposed scheme.

In contrast, UWA [16] was trained for 70 epochs by using stochastic gradient descent (SGD) with momentum for 70 epochs. The initial learning rate was set to 0.1, and was lowered by a factor of 10 at 40 and 60 epochs. We utilized a weight decay of 0.0005, a momentum of 0.9, and a batch size of 128.

In all experiments, we did not carry out any data augmentation and pre-processing to training images.

2) *Results*: Examples of reconstructed images are shown in Fig. 7, where Fig. 1(a) is the original one. Although visual information of images encrypted by the conventional methods was reconstructed by GA, it was difficult to reconstruct the visual information on plain images from images encrypted by the proposed scheme.

In addition, we measured average structural similarity (SSIM) values between original images and reconstructed ones, where lower values mean lower visual information. As shown in Table I, the proposed scheme provided the lowest SSIM values in encryption methods, namely, images generated by the proposed scheme have higher robustness against GA than conventional ones.

##### D. Image Classification

1) *Training Conditions*: We used deep residual networks (ResNet) [28] without pre-trained weights for evaluating the effectiveness of the proposed scheme in terms of image classification accuracy.

We evaluated the performance by using ResNet with 18 layers (ResNet-18). The models with ResNet were trained by using stochastic gradient descent (SGD) with momentum for 200 epochs. The initial learning rate was set to 0.1, and was lowered by a factor of 5 at 60, 120, 160 epochs. We set a weight decay and momentum to 0.0005 and 0.9.

As shown in Fig. 5(a), the models were trained with live augmentation (random cropping with padding = 4 and random horizontal flip). A client firstly carried out standard normalization and data augmentation to generate pre-processed training data ( $T'$ ) from  $T$ . Eventually, the client protected  $T'$  by using  $h_p(\cdot)$  and then uploaded  $h_p(T')$  to the servers to train DNN models. As a result, the DNN models were trained by 50K images for both CIFAR-10 and CIFAR-100.

2) *Results*: We compared the proposed scheme with the pixel-based image encryption [15] in terms of classification accuracy and robustness against COAs. As shown in Table. I, the proposed method provided almost the same classification accuracy as the pixel-based method. Therefore, the proposed scheme enables us not only to maintain the classification accuracy as that of the pixel-based method, but also to enhance robustness against DNN-based attacks.

## V. CONCLUSION

We presented a novel image transformation scheme using generative adversarial networks (GANs) for privacy-preserving deep neural networks (DNNs). The proposed scheme enables us not only to apply images without visual information to DNNs, but also to enhance robustness against ciphertext-only attack (COAs). In this paper, the proposed scheme was demonstrated to be able to protect visual information on plain images, and the visually-protected images were directly applied to DNNs for privacy-preserving image classification. Since the proposed scheme utilizes GANs, there is no need to manage encryption keys. In an experiment, we evaluated the effectiveness of the proposed scheme in terms of classification accuracy and robustness against attacks. The experimental results confirmed that the proposed scheme enables us not only to maintain high classification accuracy as that with the pixel-based method, but also to enhance robustness against DNN-based attacks.

## REFERENCES

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, p. 436, 2015.
- [2] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell, "Decaf: A deep convolutional activation feature for generic visual recognition," in *Proceedings of the 31st International Conference on Machine Learning*, Beijing, China, June 2014, pp. 647–655.
- [3] A. Krizhevsky, I. Sutskever, and G. Hinton, "Imagenet classification with deep convolutional neural networks," in *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1*, USA, 2012, pp. 1097–1105. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2999134.2999257>
- [4] A. M. Saxe, Y. Bansal, J. D. M. Advani, A. Kolchinsky, B. D. Tracey, and D. D. Cox, "On the information bottleneck theory of deep learning," in *International Conference on Learning Representations*, 2018.
- [5] M. Fredrikson, S. Jha, and T. Ristenpart, "Model inversion attacks that exploit confidence information and basic countermeasures," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 1322–1333.
- [6] S. Basu, R. Izmailov, and C. Mesterharm, "Membership Model Inversion Attacks for Deep Networks," *arXiv e-prints*, p. arXiv:1910.04257, Oct 2019.
- [7] K. Kurihara, S. Shiota, and H. Kiya, "An encryption-then-compression system for jpeg standard," in *Picture Coding Symposium (PCS)*, 2015, pp. 119–123.
- [8] M. T. Gaata and F. F. Hantoosh, "An efficient image encryption technique using chaotic logistic map and rc4 stream cipher," *International Journal of Modern Trends in Engineering and Research*, vol. 3, no. 9, pp. 213–218, 2016.
- [9] W. Sirichotedumrong, T. Chuman, S. Imaizumi, and H. Kiya, "Grayscale-based block scrambling image encryption for social networking services," in *2018 IEEE International Conference on Multimedia and Expo (ICME)*, July 2018, pp. 1–6.
- [10] M. Tanaka, "Learnable image encryption," in *2018 IEEE International Conference on Consumer Electronics-Taiwan (ICCE-TW)*, May 2018, pp. 1–2.
- [11] V. Itier, P. Puteaux, and W. Puech, "Recompression of jpeg crypto-compressed images without a key," *IEEE Transactions on Circuits and Systems for Video Technology*, pp. 1–1, 2019.
- [12] T. Chuman, W. Sirichotedumrong, and H. Kiya, "Encryption-then-compression systems using grayscale-based image encryption for jpeg images," *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 6, pp. 1515–1525, June 2019.
- [13] W. Sirichotedumrong and H. Kiya, "Grayscale-based block scrambling image encryption using ycbcr color space for encryption-then-compression systems," *APSIPA Transactions on Signal and Information Processing*, vol. 8, p. e7, 2019.
- [14] W. Sirichotedumrong, T. Maekawa, Y. Kinoshita, and H. Kiya, "Privacy-preserving deep neural networks with pixel-based image encryption considering data augmentation in the encrypted domain," in *2019 IEEE International Conference on Image Processing (ICIP)*, Sep. 2019, pp. 674–678.
- [15] W. Sirichotedumrong, Y. Kinoshita, and H. Kiya, "Pixel-based image encryption without key management for privacy-preserving deep neural networks," *IEEE Access*, vol. 7, pp. 177 844–177 855, 2019.
- [16] W. Sirichotedumrong, Y. Kinoshita, and H. Kiya, "On the security of pixel-based image encryption for privacy-preserving deep neural networks," in *2019 IEEE 8th Global Conference on Consumer Electronics (GCCE)*, October 2019, pp. 123–126.
- [17] K. Madono, M. Tanaka, M. Onishi, and T. Ogawa, "An adversarial attack to learnable encrypted images," in *22nd IEICE Symposium on Image Recognition and Understanding*, July 2019.
- [18] T. Maekawa, A. Kawamura, Y. Kinoshita, and H. Kiya, "Privacy-preserving svm computing in the encrypted domain," in *Proceedings of APSIPA Annual Summit and Conference*, 2018, pp. 897–902.
- [19] T. Maekawa, A. Kawamura, T. Nakachi, and H. Kiya, "Privacy-preserving support vector machine computing using random unitary transformation," *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, vol. E102.A, no. 12, pp. 1849–1855, 2019.
- [20] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," *arXiv:1706.06083*, 2017. [Online]. Available: <https://arxiv.org/abs/1706.06083>.
- [21] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2223–2232.
- [22] J. Johnson and A. Alahi and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *European Conference on Computer Vision*, 2016.
- [23] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *International Conference on Learning Representations*, 2015.
- [24] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention*, 2015, pp. 234–241.
- [25] A. Krizhevsky, "Learning multiple layers of features from tiny images," Tech. Rep., 2009.
- [26] A. Coates, A. Ng, and H. Lee, "An analysis of single-layer networks in unsupervised feature learning," in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, vol. 15, 11–13 Apr 2011, pp. 215–223.
- [27] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations (ICLR)*, 2015. [Online]. Available: <http://arxiv.org/abs/1412.6980>
- [28] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 770–778.