

Received February 11, 2022, accepted February 25, 2022, date of publication March 8, 2022, date of current version March 29, 2022. *Digital Object Identifier* 10.1109/ACCESS.2022.3157297

Robust Attentive Deep Neural Network for Detecting GAN-Generated Faces

HUI GUO^{®1}, SHU HU^{®2}, XIN WANG^{2,3}, (Senior Member, IEEE), MING-CHING CHANG^{®1}, (Senior Member, IEEE), AND SIWEI LYU^{®2}, (Fellow, IEEE)

¹Department of Computer Science, College of Engineering and Applied Sciences, University at Albany, State University of New York, Albany, NY 12222, USA ²Department of Computer Science and Engineering, University at Buffalo, State University of New York, Buffalo, NY 14260, USA ³Keya Medical, Seattle, WA 98104, USA

Corresponding author: Ming-Ching Chang (mchang2@albany.edu)

This work was supported in part by the U.S. Defense Advanced Research Projects Agency (DARPA) Semantic Forensic (SemaFor) Program under Kitware, Inc., through the Project Semantic Information Defender (SID) for the State University of New York and Albany and Buffalo under Grant K003088-00-S04.

ABSTRACT Generative Adversarial Network (GAN) based techniques can generate and synthesize realistic faces that cause profound social concerns and security problems. Existing methods for detecting GAN-generated faces can perform well on limited public datasets. However, images from existing datasets do not represent real-world scenarios well enough in terms of view variations and data distributions, where real faces largely outnumber synthetic ones. The state-of-the-art methods do not generalize well in real-world problems and lack the interpretability of detection results. Performance of existing GAN-face detection models degrades accordingly when facing data imbalance issues. To address these shortcomings, we propose a robust, attentive, end-to-end framework that spots GAN-generated faces by analyzing eye inconsistencies. Our model automatically learns to identify inconsistent eye components by localizing and comparing artifacts between eyes. After the iris regions are extracted by Mask-RCNN, we design a *Residual Attention Network* (*RAN*) to examine the consistency between the *corneal specular highlights* of the two eyes. Our method can effectively learn from imbalanced data using a joint loss function combining the traditional cross-entropy loss with a relaxation of the ROC-AUC loss via Wilcoxon-Mann-Whitney (WMW) statistics. Comprehensive evaluations on a newly created **FFHQ-GAN** dataset in both balanced and imbalanced scenarios demonstrate the superiority of our method.

INDEX TERMS GAN-generated face, fake face detection, iris detection, corneal specular highlights, residual attention network, data imbalance, AUC maximization, WMW statistics, FFHQ-GAN dataset.

I. INTRODUCTION

The development of Generative Adversarial Networks (GANs) [1] has led to a dramatic increase in the realism in generating high-quality face images, including PGGAN [2], StyleGAN [3], StyleGAN2 [4], and StyleGAN3 [5]. As illustrated in Figure 1, these GAN generated (or synthesized) fake faces are difficult to distinguish from human eyes. Such synthesized faces are easily generatable, can be directly leveraged for disinformation, and potentially lead to profound social, security, and ethical concerns. The GAN-generated faces can be easily abused for malicious purposes, such as creating fake social media accounts to lure or deceive

The associate editor coordinating the review of this manuscript and approving it for publication was Sudhakar Radhakrishnan¹⁰.

unaware users [6]–[9], which can cause significant security problems and frauds. Therefore, the authentication of GAN-generated faces has obtained increasing importance in recent years. However, there exists only a paucity of forensic techniques that can effectively detect such fake faces.

Many studies employ CNNs or other classifiers to distinguish the GAN-generated faces from the real ones [10]–[15]. Although these methods detect various GAN-generated faces with relatively high accuracy, similar to other deep learningbased techniques, they suffer from poor generalization and lack interpretability of detection results. Physiology-based methods [10], [16], [17] detect fake faces by examining the semantic aspects of human faces, including physiological or shape-related cues such as symmetry, iris color, and pupil shapes. Our prior work of an explainable physical method



FIGURE 1. StyleGAN2 [4] generated faces are highly realistic and can be easily abused for malicious purposes. Effective forensic methods for identifying them is of strong needs.

in [18] addressed some of the above limitations, where GAN-generated faces are identified based on a rule-based decision over the inconsistencies of the specular eye patterns. However, this method relies on assumptions of a frontal face as input and the existence of far-away lighting reflection source(s) from both eyes. When these assumptions are violated, generalization will be limited and false positives may rise significantly.

In this paper, we improve our prior method of [18] and develop an end-to-end approach for detecting GAN-generated faces by examining the inconsistencies between the corneal specular highlights of the two eyes. We first use Mask R-CNN [19] to detect and localize the iris regions. Instead of segmenting the corneal specular highlights using low-level image processing methods in [18], we design a **Residual Attention Network** which consists of *residual attention blocks* inspired from [20], to automatically learn to localize the inconsistencies. Our new method is data-driven and can better spot inconsistent artifacts, including but not limited to the corneal specular highlights.

Data imbalance is an important issue that is less addressed in existing GAN-generated face detection works. In realworld use scenarios of face examination, real face images usually outnumber GAN-generated ones by a large amount. Imbalanced data lead to learning problems and thus affect model design. It is well-known the widely-used cross-entropy loss [21] is not suitable for classifying imbalanced data. Although substantial progress is made by sampling [22], adjusting of class weights, data enhancement [23], *etc.*, learning with imbalanced data is still challenging. It is intuitive that the Area Under Curve (AUC) of the Receiver Operating Characteristic (ROC) plot can be incorporated as a loss term to improve classification performance [24]. However, the AUC is a pairwise rank-based metric with discontinuous values among iterations. Therefore, AUC is not directly applicable for loss design for end-to-end optimization of the classifier. To this end, we incorporate a ROC-AUC loss term by maximizing the Wilcoxon-Mann-Whitney (WMW) statistics of the ROC, which is shown to provide similar effects in approximating the AUC optimization [25], [26]. Our experimental results show that a combination of the binary cross-entropy loss and the WMW-AUC loss leads to the best end-to-end result.

We perform experiments on two data sources: (1) real human face images obtained from the Flickr-Faces-HQ (FFHQ) dataset [3] and (2) GAN-generated face images available from http://thispersondoesnotexist.com as shown in Figure 1. Experiment results demonstrate the superiority of the proposed method in distinguishing GAN-generated faces from the real ones. We summarize the main contributions of this paper in the following:

- We propose an end-to-end method for detecting GAN-generated faces by visually comparing the two eyes. A residual attention network model is incorporated to better focus on the inconsistencies of the eyes *e.g.* corneal specular highlights and other artifacts. Our fake face detection method is interpretable, and the proposed cues can be leveraged by human beings as well to perform examinations.
- We introduce the WMW-AUC loss that approximates the direct optimization of the AUC. This can also effectively address the data imbalance learning problem in contrast to other sampling or data augmentation approaches.
- We generate a new **FFHQ-GAN** dataset by combining portions of the FFHQ real faces with the StyleGAN2 generated images. The performance of GAN-generated face detection is evaluated on this FFHQ-GAN dataset for both balanced and imbalanced data conditions. Experimental results show that our method achieves plausible performance, especially on imbalanced datasets. The ablation study also validates the effectiveness of the proposed attention module and the loss design.

The paper is organized as follows. Section II summarizes related works on GAN-generated faces detection, attention methods, and learning from imbalanced data. Section III describes the proposed network architecture and proposed loss terms for robust learning. Section IV shows experimental results with qualitative visualization and quantitative analysis. Section V concludes this work.

II. RELATED WORKS

We briefly review related works, including GAN-generated face detection methods. We also review the literature on the attention mechanism and learning from imbalanced data.



FIGURE 2. The proposed architecture for GAN-generated face detection. We first use DLib [27] to detect faces and localize eyes, and use Mask R-CNN [19] to segment out the iris regions. A *Residual Attention Network (RAN)* then performs binary classification on the extracted *iris pair* to determine if the face is real or fake. The training is carried out using a joint loss combining the *Binary Cross-Entropy (BCE) loss* and the *ROC-AUC loss with WMW relaxation* to better handle the learning from imbalanced data (see text).

A. GAN-GENERATED FACE DETECTION

GAN-generated faces detection methods can be organized into two categories.

Data-driven methods [28]–[32] mostly train a deep neural network model to distinguish real and GAN-generated faces. These deep learning (DL) based methods work well in many scenarios, as they can better learn representations in a high-dimensional feature space instead of raw image pixels.

Physical and physiological methods look for signal traces, artifacts, or inconsistencies left by the GAN synthesizers. These methods are explainable in nature. Simple cues such as color difference are used in [33], [34] to distinguish GAN images from the real ones. However, those methods are no longer effective as the GAN methods advance. More sophisticated methods in [11], [35] leverage fingerprints or abstract signal-level traces of the noise residuals to differentiate GAN-generated faces. Many works [36]-[38] identify GAN images by recognizing the specific artifacts produced by the GAN upsampling process. In [39], the distribution of facial landmarks is analyzed to distinguish GAN-generated faces. Inconsistent head poses are detected to expose the fake videos in [10]. The work of [16] identifies GAN-generated faces as well as deepfakes face manipulations by inspecting visual artifacts. Our prior work of [18] determines the inconsistencies of the corneal specular highlights between left and right eyes to expose GAN-generated faces.

B. ATTENTION MECHANISM

Since the seminal work of [40] in machine translation, the *attention mechanism* is widely used in many applications on improving the performance of deep learning models by focusing on the most relevant part of the features in a flexible manner. The Class Activation Mapping (CAM) [41] and Grad-CAM [42] are widely used in many computer vision tasks [43]. However, in these works, attentions are only used to visualize model prediction in showing significant portions of the images. On the other hand, integrating the attention mechanism into the network design is shown to be effective in boosting performance, as the network can be guided by the attention to focus on relevant regions during training [44].

The channel attention [45] can automatically learn to focus on important channels by analyzing the relationship between channels. SENet [46] embeds the channel attention mechanism into residual blocks, and effectiveness is shown on large-scale image classification. The attention mechanism is also used in [47] to distinguish important channels in the network to improve the representation capability. The idea of channel attention and spatial attention are combined jointly in [48], [49] to improve network performance significantly. The Residual Attention Network in [20] combines the residual unit [50] with the attention mechanism by stacking residual attention blocks to improve performance and reduce model complexity.

C. IMBALANCED DATA LEARNING

Learning from imbalanced data has been widely studied in machine learning [51]–[54] and computer vision [55], [56]. Earlier solutions for imbalanced data learning are mainly based on the sampling design, *e.g. oversampling* for minor classes, *undersampling* for major classes, and weighed sampling [57], *etc.* These sampling-based methods come with their own drawbacks. For example, undersampling may ignore important samples, and oversampling may lead to overfitting.

Data augmentation provides an alternative solution to alleviate data imbalance issues. For image recognition, image mirroring, rotation, color adjustments, *etc.* are simple methods to augment data samples [58]. However, data augmentation methods can only address the data imbalance problems partly, as the size of the original dataset must be diverse enough, such that a sufficient amount of representative samples can be produced from augmentation.

III. METHOD

We next describe the proposed GAN-generated faces detection framework. Given an input face image, facial landmarks are first localized using DLib [27], and Mask R-CNN [19] is used to segment out the left and right iris regions of the eyes (§ III-A). We adopt a residual attention-based network [20] to perform binary classification on the iris regions of interest to determine if the input image is real or fake (§ III-B).



FIGURE 3. Details of our Attention Module from the our RAN in Figure 2. Our design is inspired from the residual attention network of [20].

The training of our network aims to maximize the classification performance reflected by the standard Area-Under-Curve of the ROC plot (§ III-C), which is general and can effectively address the data imbalance problem. However, due to the discrete nature of the ROC-AUC values, a naive gradient-based implementation does not work for end-to-end learning. In § III-D), we present a detailed solution in our proposed approach by relaxing the AUC maximization and approximating the goal using the WMW statistics. Figure 5 overviews the training pipeline of the proposed method that can effectively learn from imbalanced data.

A. FACIAL LANDMARK LOCALIZATION AND IRIS SEGMENTATION

Given a face image, the first step of our method to determine if it is real or GAN-generated is to detect and localize the face using the facial landmark extractor provided in DLib [27]. The localized regions containing the eyes are cropped out for consistency checking. Mask R-CNN [19], the state-ofthe-art detection and segmentation network, is employed to further detect and localize the iris regions. Mask R-CNN is a two-stage network based on Faster R-CNN [59] as shown in Figure 2 (middle). The first stage of Mask R-CNN is a Region Proposal Network (RPN) that generates candidate object bounding boxes for all the object categories. In the second stage, the R-CNN extracts features using the Region of Interest Align (RoIAlign) layer for each proposal. In the last stage, label classification and bounding box regression are performed for each proposal, and mask prediction is performed in a parallel branch. We train the Mask R-CNN model using the eye region images from the datasets in [60], [61], where more details will be provided in § IV-B. Figure 4 shows examples of the extracted pair of iris regions from the cases of GAN-generated (left) and real (right) faces.

B. RESIDUAL ATTENTION NETWORK

We adopt the attention mechanism [45], [47], [62]–[64] to improve the spotting of inconsistent corneal specular

TABLE 1.	Details of the propose	ed Residual	Attention	Network	(RAN) in
the right o	of Figure 2.				

Layer	Output Size	Network
Conv1	96×96×16	3×3 , stride 1
Max pooling	48×48×16	3×3 , stride 2
Residual block	48×48×64	$1 \times \begin{bmatrix} 1 \times 1, 16 \\ 3 \times 3, 16 \\ 1 \times 1, 64 \\ 1 \times 1, 64 \end{bmatrix}$
Attention Module	$48 \times 48 \times 64$	$2 \times$ attention
Residual block	24×24×128	$1 \times \begin{bmatrix} 1 \times 1, 32 \\ 3 \times 3, 32 \\ 1 \times 1, 128 \\ 1 \times 1, 128 \end{bmatrix}$
Attention Module	$24 \times 24 \times 128$	$2 \times$ attention
Residual block	12×12×256	$1 \times \begin{bmatrix} 1 \times 1, 64 \\ 3 \times 3, 64 \\ 1 \times 1, 256 \\ 1 \times 1, 256 \end{bmatrix}$
Attention Module	12×12×256	$2 \times$ attention
Residual block	6×6×512	$3 \times \begin{bmatrix} 1 \times 1, 128 \\ 3 \times 3, 128 \\ 1 \times 1, 512 \\ 1 \times 1, 512 \end{bmatrix}$
Average pooling	1×1×512	6×6 , stride 1
FC, Sigmoid	1	



(a) from **GAN-generated** faces

(b) from real faces.

FIGURE 4. The extracted iris pairs of our method for the (a) GAN-generated and (b) real faces. Artifacts of inconsistent corneal specular highlights are obvious in GAN-generated iris pairs.

highlights between the eyes so as to improve GAN-generated face detection. Incorporating attention to a detection/ segmentation network is commonly accomplished by having a separate branch that calculates the attention maps and later is incorporated back to the main branch with weights. Inspired from [20], each Attention Module in our attention network consists of a *trunk branch* and a *soft mask branch*. The trunk branch contains several residual blocks [50] and acts as a shortcut for data flow. The soft mask branch uses a U-net structure [65] to weight output features. Specifically, given the input feature map f, denote the output of the trunk branch T as T(f), and the output of the soft mask branch M as M(f), respectively. As illustrated in Figure 3, the final attended feature map f' is obtained via element-wise matrix product via

$$f' = (1 + M(f)) \circ T(f),$$
 (1)

where the symbol \circ denotes the Hadamard product.

The attention module can be configured to focus learning on channel attention, spatial attention, or mixed attention.



FIGURE 5. The proposed pipeline for training the Residual Attention Network (RAN) on possibly imbalanced data for GAN-generated face classification. The extracted iris pairs are passed as input to the RAN. A robust loss function derived from maximizing the AUC of ROC is optimized in the training of the RAN. See details in § III-D.

As suggested in [20], the mixed attention yields the best performance. Thus, we use the Sigmoid function $\frac{1}{1+\exp(-f_{s,c})}$ to learn the mixed attention for each channel and each spatial location, where s ranges over all spatial positions and cranges over all channels of f. The proposed Residual Attention Network (RAN) is constructed by stacking multiple Attention Modules, as shown on the right side of Figure 2. Table 1 provides details of the architectures. Although the attention module plays an important role in classification, a simple stacking of attention modules may reduce performance. To this end, we adopt a simple solution by adding the attention map onto the original feature map. This combination allows attention modules stacked like a ResNet [50] and improves the performance [20]. Given an input image, The RAN outputs a prediction score from the last Sigmoid layer, as an indication of the likelihood of the input image being a GAN-generated image.

C. AUC OF ROC FOR CLASSIFICATION EVALUATION

Most classification loss measures including the popular cross-entropy loss are ineffective in addressing the issue of data imbalance. The resulting models can produce accurate but rather biased predictions that do not work well in practice. It is desirable to address data imbalance directly by specifically designing a suitable loss function.

Since the area under the curve (AUC) of a receiver operation curve (ROC) [26], [66] is a robust evaluation metric for both balanced and imbalanced data, we would like to directly maximize the AUC to handle imbalanced situations. The AUC is widely used in the binary classification problems. We next briefly review the definition of AUC, and then motivate how we incorporate a loss term that directly maximize the AUC performance. Given a labeled dataset $\{(\mathbf{x}_i, y_i)\}_{i=1}^M$, where each data sample $\mathbf{x}_i \in \mathbb{R}^d$ and each corresponding label $y_i \in \{-1, +1\}$. We define a set of indices of positive instances as $\mathcal{P} = \{i \mid y_i = +1\}$. Similarly, the set of indices of negative instances is $\mathcal{N} = \{i \mid y_i = -1\}$. Let $g_{\mathbf{w}} : \mathbb{R}^d \to \mathbb{R}$ be a parametric prediction function with parameter $\mathbf{w} \in \mathbb{R}^m$. $g_{\mathbf{w}}(\mathbf{x}_i)$ represents the prediction score of the *i*-th sample, where $i \in \{1, \dots, M\}$. For simplicity, we assume $g_{\mathbf{w}}(\mathbf{x}_i) \neq g_{\mathbf{w}}(\mathbf{x}_j)$ for $i \neq j$ (ties can be broken in any consistent way).

Given a threshold λ , the number of negative examples with prediction scores larger than λ is false positive (FP), and the number of positive examples with prediction scores greater or equal to λ is true positive (TP). According to the FP and TP, we can calculate the false positive rate (FPR) and the true positive rate (TPR) as follows,

$$FPR = \frac{\sum_{i \in \mathcal{N}} \mathbb{I}_{[g_{\mathbf{w}}(\mathbf{x}_i) > \lambda]}}{|\mathcal{N}|}, \quad TPR = \frac{\sum_{i \in \mathcal{P}} \mathbb{I}_{[g_{\mathbf{w}}(\mathbf{x}_i) \geq \lambda]}}{|\mathcal{P}|},$$

where $\mathbb{I}_{[a]}$ is an indicator function with $\mathbb{I}_{[a]} = 1$ if *a* is true and 0 otherwise. The receiver operation curve (ROC) is a plot of FPR versus TPR with setting different decision thresholds $\lambda \in (-\infty, \infty)$. Based on this definition, ROC is a curve confined to $[0, 1] \times [0, 1]$ and connecting the point (0,0) to (1,1). The value of AUC corresponds to the area enclosed by the ROC curve.

D. WMW AUC RELAXATION FOR LOSS DESIGN

The computation of an AUC score based on the area under a ROC curve cannot be directly used in a loss function due to its discrete nature. According to the Wilcoxon-Mann-Whitney (WMW) statistic [25], we can relax the AUC as follows,

$$AUC = \frac{1}{|\mathcal{P}||\mathcal{N}|} \sum_{i \in \mathcal{P}} \sum_{j \in \mathcal{N}} \mathbb{I}_{[g_{\mathbf{w}}(\mathbf{x}_i) > g_{\mathbf{w}}(\mathbf{x}_j)]}$$

Therefore, the corresponding AUC loss (risk) can be defined as:

$$\mathcal{L}_{AUC} = 1 - \text{AUC} = \frac{1}{|\mathcal{P}||\mathcal{N}|} \sum_{i \in \mathcal{P}} \sum_{j \in \mathcal{N}} \mathbb{I}_{[g_{\mathbf{w}}(\mathbf{x}_i) < g_{\mathbf{w}}(\mathbf{x}_j)]}.$$
 (2)

Obviously, \mathcal{L}_{AUC} takes value in [0, 1]. It is a fraction of pairs of prediction scores from the positive sample and negative sample that are ranked incorrectly, *i.e.*, the prediction score from a negative sample is larger than the prediction score from a positive sample. If all prediction scores from the positive samples are larger than any prediction score from the negative samples, then $\mathcal{L}_{AUC} = 0$. This indicates we obtain a perfect classifier. Furthermore, \mathcal{L}_{AUC} is independent

of the threshold λ . \mathcal{L}_{AUC} only depends on the prediction scores $g_{\mathbf{w}}(\mathbf{x})$. In other words, the predictor $g_{\mathbf{w}}$ affects the value of \mathcal{L}_{AUC} . Therefore, we aim to learn a classifier $g_{\mathbf{w}}$ that minimizes Eq.(2).

Although we can calculate \mathcal{L}_{AUC} by comparing prediction score from the positive sample and prediction score from the negative sample in each pair, the \mathcal{L}_{AUC} formulation is non-differentiable due to the discrete computation. It is therefore desirable to find a differentiable approximation for \mathcal{L}_{AUC} . Inspired by the work in [25], we find an approximation to \mathcal{L}_{AUC} that can be directly applied to our objective function to minimize the AUC loss along with our imbalanced training procedure. Specifically, a differentiable approximation of \mathcal{L}_{AUC} can be reformulated as:

$$\mathcal{L}_{AUC} = \frac{1}{|\mathcal{P}||\mathcal{N}|} \sum_{i \in \mathcal{P}} \sum_{j \in \mathcal{N}} R(g_{\mathbf{w}}(\mathbf{x}_i), g_{\mathbf{w}}(\mathbf{x}_j)), \qquad (3)$$

and $R(g_{\mathbf{w}}(\mathbf{x}_i), g_{\mathbf{w}}(\mathbf{x}_j)) =$

$$\begin{cases} (-(g_{\mathbf{w}}(\mathbf{x}_i) - g_{\mathbf{w}}(\mathbf{x}_j) - \gamma))^p, & g_{\mathbf{w}}(\mathbf{x}_i) - g_{\mathbf{w}}(\mathbf{x}_j) < \gamma, \\ 0, & \text{otherwise,} \end{cases}$$
(4)

where $\gamma \in (0, 1]$ and p > 1 are two hyperparameters.

Loss for the Proposed Residual Attention Network: We use a joint loss function comprising the conventional binary cross-entropy (BCE) loss function \mathcal{L}_{BCE} and the AUC loss function \mathcal{L}_{AUC} in weighted sum:

$$\mathcal{L} = \alpha \, \mathcal{L}_{BCE} + (1 - \alpha) \, \mathcal{L}_{AUC}, \qquad (5)$$

where $\alpha \in [0, 1]$ is a scaling factor that is designed for balancing the weights of the BCE loss and the AUC loss.

IV. EXPERIMENT

For experimental evaluation of the proposed method and comparison against the state-of-the-art methods, we first introduce the newly constructed FFHQ-GAN datasets § IV-A. Implementation details of the proposed method are provided in § IV-B. Performance evaluation on the FFHQ-GAN balanced and imbalanced subsets is in § IV-C. Ablation studies are provided in § IV-D. Finally, qualitative results are shown in § IV-E.

A. THE NEW FFHQ-GAN DATASET

We collect real human face images from the Flickr-Faces-HQ (FFHQ) dataset [3]. GAN-generated face images are created using StyleGAN2 [3] via http://thispersondoesnotexist.com, where the image resolution is 1024×1024 pixels. We randomly select 5,000 real face images from FFHQ and 5,000 GAN-generated face images. After iris detection, we discard those images with the iris of any eye not detected. This ends up with 3,739 real faces (with iris pairs) and 3,748 fake faces (with iris pairs), which constitutes our new **FFHQ-GAN dataset**. The split ratio of training and testing is 8:2.

To enable a thorough evaluation of the model in both balanced and imbalanced data scenarios, we sampled the FFHQ-GAN dataset to form an imbalanced subset, where the statistics of the subsets are provided in Table 2.
 TABLE 2. Details of the FFHQ-GAN dataset regarding its balanced (-b)

 and imbalanced (-imb) subsets.

Detecate	Training Set		Testing Set		Datio ~	
Datasets	GAN-face	Real	GAN-face	Real	$\kappa a n o \approx$	
FFHQ-GAN-b	2998	2991	750	748	1:1	
FFHQ-GAN-imb	400	2000	100	500	1:5	

TABLE 3. Results on the FFHQ-GAN dataset regarding its balanced (-b) and imbalanced (-imb) subsets.

Datasats	methods	Metric				
Datasets	methods	ACC	P	R	F1	AUC
	ResNet	0.85	0.85	0.84	0.85	0.92
FEHO CAN b	Xception	0.95	0.95	0.96	0.95	0.99
rrnQ-GAN-0	RAN with BCE loss	0.92	0.92	0.92	0.92	0.98
	RAN with BCE+AUC loss	0.97	0.98	0.97	0.97	1.00
	ResNet	0.72	0.28	0.41	0.33	0.64
FEUO CAN imh	Xception	0.87	0.75	0.36	0.49	0.85
TTHQ-OAN-IIID	RAN with BCE loss	0.89	0.84	0.41	0.55	0.89
	RAN with BCE+AUC loss	0.96	0.89	0.85	0.87	0.98

B. IMPLEMENTATION DETAILS

We implemented our method in PyTorch [67]. Experiments are conducted on a workstation with two NVIDIA GeForce 1080Ti GPUs.

For iris detection, Mask R-CNN is trained using the datasets from [60], [61]. For each training eye image, the outer boundary mask of each iris is obtained using the method of [60] with default hyper-parameter settings. These masks are used to generate the iris bounding boxes and the corresponding masks for training, using the default settings in [19]. In the test stage, given an input face image, we first use the face detector and landmark extractor of DLib [27] to crop out the eye regions. Each cropped eye region is fed to Mask R-CNN for localizing the iris bounding box and segmentation mask. This process is repeated for both the left and right eyes to obtain the iris pairs as the input for our Residual Attention Network. We resize all iris pairs to a fixed size 96×96 for training and testing to ensure that the whole pipeline works well.

Table 1 describes the details of our Residual Attention Network (RAN), where the Attention Module (AM) detailed in Figure 3 is repeatedly stacked three times. The network is trained using Adam optimizer [68] with the learning rate of 0.001 and batch size 128. Training is terminated at 100 epochs for balanced data and 2,000 for imbalanced data.

Hyper-Parameters: We set p = 2 in Eq. (4) and $\gamma = 0.4$ for balanced dataset, and $\gamma = 0.6$ for imbalanced data. For the experiments on the balanced dataset, α in Eq. (5) is set to 0.2. For the experiments on the imbalanced dataset, α is set to 0.4. These hyperparameters yields the best performance.

C. EVALUATION ON THE FFHQ-GAN DATASET

We report evaluation of GAN-generated face detection on the FFHQ-GAN dataset in terms of Accuracy (ACC), Precision (P), Recall (R), F1 score (F1), the area under the curve (AUC)



FIGURE 6. Performance comparison of the proposed method with ResNet with BCE loss, Xception with BCE loss, and RAN with BCE loss.



FIGURE 7. Confusion Matrix on the FFHQ-GAN (left) balanced and (right) imbalanced datasets.



FIGURE 8. Impact of hyperparameter α of the AUC loss in Eq.(4) for the GAN-generated face detection on the imbalanced dataset.

of the ROC, and Precision-Recall (PR) curves. Accuracy is calculated as ACC = $\frac{TP+TN}{TP+TN+FP+FN}$, where FN and TN indicate false negatives and true negatives, respectively. Precision-Recall is calculated as P = $\frac{TP}{TP+FP}$ and R = $\frac{TP}{TP+FN}$. F1 score is the harmonic average value of P and R, as F1 = $\frac{2PR}{P+R}$.

Results on the Balanced and Imbalanced Set: To evaluate the effectiveness of the proposed method, we evaluate RAN trained with BCE + AUC loss against the ResNet-50 [50] and Xception [69], two of the widely-used DNN classification models trained with the BCE loss. Table 3 presents the classification results of experiments training/test on the balanced and imbalanced FFHQ-GAN datasets. The corresponding ROC and Precision-Recall curves are shown in Figure 6 (right), and the Confusion Matrices are shown in Figure 7. Results show that both the ResNet-50 and Xception obtain low Recall scores due to the imbalanced data distribution

(a) Iris pairs from GAN-generated faces



(b) Iris pairs from real faces



FIGURE 9. Visualization of the extracted iris pairs and the corresponding attention maps obtained from our Residual Attention Network (RAN). Observe that the attention maps for GAN-generated faces better focus on the artifacts such as the corneal specular highlights, while the attention maps for real faces are widely distributed. This shows the effective learning of RAN for identifying GAN-generated faces.

in Table 3. In comparison, our method achieves the highest performance in all metrics. These results indicate that our method can effectively improve performance on both balanced and imbalanced data training. We have also performed experiments to train our model, ResNet-50, and Xception



FIGURE 10. Examples of detected GAN-generated faces and their corresponding iris regions and the attention maps produced from our method. These examples show that our method can detect a wide range of face images, including those with tilted or side views where both irises are visible.

on a balanced dataset and test on an imbalanced dataset. The obtained performance difference is similar to that of training/test on the balanced dataset. This result suggests the importance of the training model on an imbalanced dataset if the model is expected to deal with detection on imbalanced data.

D. ABLATION STUDIES

1) EFFECT OF THE AUC LOSS

We compare the proposed RAN model trained with the ideal case with combined AUC and BCE loss in Eq. (5) against the same model trained only with BCE loss. Results are shown in Table 3 and Figures 6 and 7. Observe that the proposed joint BCE+AUC loss outperforms the same model trained only with BCE loss alone in all evaluation metrics. In other words, the incorporation of the AUC loss improves the classification performance substantially and consistently.

2) HYPER-PARAMETER ANALYSIS

We also study the impact of hyper-parameter α in our loss function in Eq. (5) regarding detection performance of imbalanced data. Figure 8 shows the experimental results of the obtained AUC score versus α ranging from 0 to 1, and $\alpha = 0.4$ yields the best detection performance.

E. QUALITATIVE RESULTS

Figure 9 provides visualization of the attention maps of the real and GAN-generated iris examples. Observe that there is an obvious difference between the corresponding attention maps of the GAN-generated irises and the real ones. Concretely, the network attends on the whole iris part for the real

images and attends to the highlight parts for the fake images. Figure 10 shows additional examples of the GAN-generated face with the extracted iris pairs and corresponding attention maps. The visualization also provides an intuitive approach for human beings to identify GAN-generated faces by comparing their iris regions.

V. CONCLUSION

In this work, we investigate building a robust end-to-end deep learning framework for detecting GAN-generated faces. We show that GAN-generated faces can be distinguished from real faces by examining the consistency between the two iris regions. In particular, artifacts such as the corneal specular highlight inconsistencies can be robustly identified through end-to-end learning via the proposed Residual Attention Network. Our design of a joint loss combining the AUC loss with the cross-entropy loss can effectively deal with the learning from imbalanced data. We also showed that a direct optimization of the ROC-AUC loss is computationally not feasible, however relaxing the ROC AUC via the Wilcoxon-Mann-Whitney (WMW) statistics can provide a good approximation. Our GAN- face detection result is explainable, and the approach of spotting iris inconsistency can also serve as a useful cue for human users. Experimental results show that our model achieves superior performance on both balanced and imbalanced datasets for GAN-generated faces detection.

REFERENCES

 I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Proc. NIPS*, 2014, pp. 1–9.

- [2] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of GANs for improved quality, stability, and variation," 2017, arXiv:1710.10196.
- [3] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 4401–4410.
- [4] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of StyleGAN," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 8110–8119.
- [5] T. Karras, M. Aittala, S. Laine, E. Härkönen, J. Hellsten, J. Lehtinen, and T. Aila, "Alias-free generative adversarial networks," in *Proc. NeurIPS*, 2021, pp. 1–12.
- [6] A Spy Reportedly Used an AI-Generated Profile Picture to Connect With Sources on LinkedIn. Accessed: Jun. 13, 2019. [Online]. Available: https://bit.ly/35BU215
- [7] A High School Student Created a Fake 2020 US Candidate. Twitter Verified it. Accessed: Feb. 28, 2020. [Online]. Available: https://www. cnn.com/2020/02/28/tech/fake-twitter-candidate-2020/index.html
- [8] How Fake Faces are Being Weaponized Online. Accessed: Feb. 20, 2020.
 [Online]. Available: https://www.cnn.com/2020/02/20/tech/fake-facesdeepfake/index.html
- [9] These Faces are not Real. Accessed: Jul. 15, 2020. [Online]. Available: https://graphics.reuters.com/CYBER-DEEPFAKE/ACTIVIST/ nmovajgnxpa/index.html
- [10] X. Yang, Y. Li, and S. Lyu, "Exposing deep fakes using inconsistent head poses," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2019, pp. 8261–8265.
- [11] F. Marra, D. Gragnaniello, L. Verdoliva, and G. Poggi, "Do GANs leave artificial fingerprints?" in *Proc. IEEE Conf. Multimedia Inf. Process. Retr.* (*MIPR*), Mar. 2019, pp. 506–511.
- [12] H. Mo, B. Chen, and W. Luo, "Fake faces identification via convolutional neural network," in *Proc. 6th ACM Workshop Inf. Hiding Multimedia Secur.*, Jun. 2018, pp. 43–47.
- [13] N.-T. Do, I.-S. Na, and S.-H. Kim, "Forensics face detection from GANs using convolutional neural network," in *Proc. ISITC*, 2018, pp. 1–5.
- [14] R. Wang, F. Juefei-Xu, L. Ma, X. Xie, Y. Huang, J. Wang, and Y. Liu, "FakeSpotter: A simple yet robust baseline for spotting AI-synthesized fake faces," 2019, arXiv:1909.06122.
- [15] B. Chen, X. Ju, B. Xiao, W. Ding, Y. Zheng, and V. H. C. de Albuquerque, "Locally GAN-generated face detection based on an improved Xception," *Inf. Sci.*, vol. 572, pp. 16–28, Sep. 2021.
- [16] F. Matern, C. Riess, and M. Stamminger, "Exploiting visual artifacts to expose deepfakes and face manipulations," in *Proc. IEEE Winter Appl. Comput. Vis. Workshops (WACVW)*, Jan. 2019, pp. 83–92.
- [17] H. Guo, S. Hu, X. Wang, M.-C. Chang, and S. Lyu, "Eyes tell all: Irregular pupil shapes reveal GAN-generated faces," in *Proc. ICASSP*, 2022, pp. 1–6.
- [18] S. Hu, Y. Li, and S. Lyu, "Exposing GAN-generated faces using inconsistent corneal specular highlights," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Jun. 2021, pp. 2500–2504.
- [19] K. He, G. Gkioxari, P. Dollar, and R. Girshick, "Mask R-CNN," in *Proc. ICCV*, 2017, pp. 2961–2969.
- [20] F. Wang, M. Jiang, C. Qian, S. Yang, C. Li, H. Zhang, X. Wang, and X. Tang, "Residual attention network for image classification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 3156–3164.
- [21] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*. Cambridge, MA, USA: MIT Press, 2012.
- [22] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2015, pp. 1–9.
- [23] M. Fadaee, A. Bisazza, and C. Monz, "Data augmentation for low-resource neural machine translation," 2017, arXiv:1705.00440.
- [24] F. Provost, T. Fawcett, and R. Kohavi, "The case against accuracy estimation for comparing induction algorithms," in *Proc. ICML*, vol. 98, 1998, pp. 445–453.
- [25] L. Yan, R. H. Dodier, M. Mozer, and R. H. Wolniewicz, "Optimizing classifier performance via an approximation to the Wilcoxon-Mann-Whitney statistic," in *Proc. ICML*, 2003, pp. 848–855.
- [26] S. Lyu and Y. Ying, "A univariate bound of area under ROC," in *Proc.* UAI, 2018, pp. 1–10.
- [27] D. E. King, "Dlib-ml: A machine learning toolkit," J. Mach. Learn. Res., vol. 10, pp. 1755–1758, Jan. 2009.

- [28] F. Marra, C. Saltori, G. Boato, and L. Verdoliva, "Incremental learning for the detection and classification of GAN-generated images," in *Proc. IEEE Int. Workshop Inf. Forensics Secur. (WIFS)*, Dec. 2019, pp. 1–6.
- [29] M. Goebel, L. Nataraj, T. Nanjundaswamy, T. M. Mohammed, S. Chandrasekaran, and B. S. Manjunath, "Detection, attribution and localization of GAN generated images," 2020, arXiv:2007.10466.
- [30] S.-Y. Wang, O. Wang, R. Zhang, A. Owens, and A. A. Efros, "CNNgenerated images are surprisingly easy to spot...for now," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, vol. 7, Jun. 2020, pp. 1–10.
- [31] Z. Liu, X. Qi, and P. H. S. Torr, "Global texture enhancement for fake face detection in the wild," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 8060–8069.
- [32] N. Hulzebosch, S. Ibrahimi, and M. Worring, "Detecting CNN-generated facial images in real-world scenarios," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2020, pp. 642–643.
- [33] S. McCloskey and M. Albright, "Detecting GAN-generated imagery using color cues," 2018, arXiv:1812.08247.
- [34] H. Li, B. Li, S. Tan, and J. Huang, "Identification of deep network generated images using disparities in color components," 2018, arXiv:1808.07276.
- [35] N. Yu, L. Davis, and M. Fritz, "Attributing fake images to GANs: Learning and analyzing GAN fingerprints," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 7556–7566.
- [36] X. Zhang, S. Karaman, and S.-F. Chang, "Detecting and simulating artifacts in GAN fake images," in *Proc. IEEE Int. Workshop Inf. Forensics Secur. (WIFS)*, Dec. 2019, pp. 1–6.
- [37] J. Frank, T. Eisenhofer, L. Schönherr, A. Fischer, D. Kolossa, and T. Holz, "Leveraging frequency analysis for deep fake image recognition," 2020, arXiv:2003.08685.
- [38] R. Durall, M. Keuper, and J. Keuper, "Watch your up-convolution: CNN based generative deep neural networks are failing to reproduce spectral distributions," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.* (CVPR), Jun. 2020, pp. 7890–7899.
- [39] X. Yang, Y. Li, H. Qi, and S. Lyu, "Exposing GAN-synthesized faces using landmark locations," in *Proc. ACM Workshop Inf. Hiding Multimedia Secur.*, Jul. 2019, pp. 113–118.
- [40] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," 2014, arXiv:1409.0473.
- [41] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2921–2929.
- [42] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Int. Conf. Comput. Vis.* (*ICCV*), Oct. 2017, pp. 618–626.
- [43] A. Chattopadhay, A. Sarkar, P. Howlader, and V. N. Balasubramanian, "Grad-CAM++: Generalized gradient-based visual explanations for deep convolutional networks," in *Proc. IEEE Winter Conf. Appl. Comput. Vis.* (WACV), Mar. 2018, pp. 839–847.
- [44] S. Kardakis, I. Perikos, F. Grivokostopoulou, and I. Hatzilygeroudis, "Examining attention mechanisms in deep learning models for sentiment analysis," *Appl. Sci.*, vol. 11, no. 9, p. 3883, Apr. 2021.
- [45] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, "CBAM: Convolutional block attention module," in *Proc. ECCV*, 2018, pp. 3–19.
- [46] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit., Jun. 2018, pp. 7132–7141.
- [47] Y. Zhang, K. Li, K. Li, L. Wang, B. Zhong, and Y. Fu, "Image superresolution using very deep residual channel attention networks," in *Proc. ECCV*, 2018, pp. 286–301.
- [48] L. Chen, H. Zhang, J. Xiao, L. Nie, J. Shao, W. Liu, and T.-S. Chua, "SCA-CNN: Spatial and channel-wise attention in convolutional networks for image captioning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* (CVPR), Jul. 2017, pp. 5659–5667.
- [49] S. Woo, J. Park, J. Y. Lee, and I. S. Kweon, "CBAM: Convolutional block attention module," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 3–19.
- [50] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [51] Y. Yang and Z. Xu, "Rethinking the value of labels for improving classimbalanced learning," 2020, arXiv:2006.07529.

- [52] K. Cao, C. Wei, A. Gaidon, N. Arechiga, and T. Ma, "Learning imbalanced datasets with label-distribution-aware margin loss," in *Proc. NeurIPS*, 2019, pp. 1567–1578.
- [53] S. Hu, Y. Ying, and S. Lyu, "Learning by minimizing the sum of ranked range," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 21013–21023.
- [54] S. Hu, Y. Ying, X. Wang, and S. Lyu, "Sum of ranked range loss for supervised learning," 2021, arXiv:2106.03300.
- [55] Y. Wang, W. Gan, J. Yang, W. Wu, and J. Yan, "Dynamic curriculum learning for imbalanced data classification," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 5017–5026.
- [56] C. Huang, Y. Li, C. C. Loy, and X. Tang, "Learning deep representation for imbalanced classification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 5375–5384.
- [57] H. He and E. A. Garcia, "Learning from imbalanced data," *IEEE Trans. Knowl. Data Eng.*, vol. 21, no. 9, pp. 1263–1284, Sep. 2009.
- [58] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," J. Big Data, vol. 6, no. 1, pp. 1–48, Dec. 2019.
- [59] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," 2015, arXiv:1506.01497.
- [60] C. Wang, J. Muhammad, Y. Wang, Z. He, and Z. Sun, "Towards complete and accurate iris segmentation using deep multi-task attention network for non-cooperative iris recognition," *IEEE Trans. Inf. Forensics Security*, vol. 15, pp. 2944–2959, 2020.
- [61] C. Wang et al., "NIR iris challenge evaluation in non-cooperative environments: Segmentation and localization," in Proc. IEEE Int. Joint Conf. Biometrics (IJCB), Aug. 2021, pp. 1–10.
- [62] L.-C. Chen, Y. Yang, J. Wang, W. Xu, and A. L. Yuille, "Attention to scale: Scale-aware semantic image segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 3640–3649.
- [63] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu, "Spatial transformer networks," in *Proc. NIPS*, vol. 28, 2015, pp. 2017–2025.
- [64] Q. Jin, Z. Meng, C. Sun, H. Cui, and R. Su, "RA-UNet: A hybrid deep attention-aware network to extract liver and tumor in CT scans," *Frontiers Bioeng. Biotechnol.*, vol. 8, p. 1471, Dec. 2020.
- [65] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proc. Med. Image Comput. Comput.-Assist. Intervent. (MICCAI).* Springer, 2015, pp. 234–241.
- [66] C. Cortes and M. Mohri, "AUC optimization vs. error rate minimization," in Proc. Adv. Neural Inf. Process. Syst., vol. 16, 2003, pp. 313–320.
- [67] A. Paszke et al., "PyTorch: An imperative style, high-performance deep learning library," in Proc. Adv. Neural Inf. Process. Syst., vol. 32, Dec. 2019, pp. 8026–8037.
- [68] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015, pp. 1–15.
- [69] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1251–1258.



HUI GUO is currently pursuing the Ph.D. degree with the University at Albany, State University of New York. Her research interests include digital media forensics, computer vision, and deep learning.



SHU HU received the M.Eng. degree in software engineering from the University of Science and Technology of China, in 2016, and the M.A. degree in mathematics from the University at Albany, State University of New York, in 2020. He is currently pursuing the Ph.D. degree in computer science with the University at Buffalo, State University of New York. His research interests include machine learning, digital media forensics, and computer vision.



XIN WANG (Senior Member, IEEE) received the Ph.D. degree in computer science from the University at Albany, State University of New York, Albany, NY, USA, in 2015. He is currently a Senior Machine Learning Scientist at Keya Medical, Seattle, WA, USA. His research interests include artificial intelligence, machine learning, reinforcement learning, medical image computing, computer vision, and media forensics.



MING-CHING CHANG (Senior Member, IEEE) received the B.S. degree in civil engineering and the M.S. degree in computer science and information engineering (CSIE) from the National Taiwan University, in 1996 and 1998, respectively, and the Ph.D. degree from the Laboratory for Engineering Man/Machine Systems (LEMS), School of Engineering, Brown University, in 2008. He was an Assistant Researcher at Mechanical Industry Research Labs, Industrial Technology Research

Institute (ITRI), Taiwan, from 1996 to 1998. From 2008 to 2016, he was a Computer Scientist at the GE Global Research Center. From 2016 to 2018, he was with the Department of Electrical and Computer Engineering. He is currently an Assistant Professor at the Department of Computer Science, College of Engineering and Applied Sciences (CEAS), University at Albany, State University of New York (SUNY). His research projects are funded by GE Global Research, IARPA, DARPA, NIJ, VA, and UAlbany. He has authored more than 100 peer-reviewed journals and conference publications, seven U.S. patents, and 15 disclosures. His research interests include video analytics, computer vision, image processing, and artificial intelligence. He is a member of ACM. He was a recipient of the IEEE Advanced Video and Signal-based Surveillance (AVSS) 2011 Best Paper Award - Runner-Up, the IEEE Workshop on the Applications of Computer Vision (WACV) 2012 Best Student Paper Award, the GE Belief - Stay Lean and Go Fast Management Award in 2015, and the IEEE Smart World NVIDIA AI City Challenge 2017 Honorary Mention Award. He serves as the Co-Chair for the Annual AI City Challenge CVPR 2018-2021 Workshop, the Co-Chair for the IEEE Lower Power Computer Vision (LPCV) Annual Contest and Workshop 2019-2021, the Program Chair for the IEEE Advanced Video and Signal-Based Surveillance (AVSS) 2019, the Co-Chair for the IWT4S 2017-2019, the Area Chair for IEEE ICIP (2017, 2019-2021) and ICME (2021), and the TPC Chair for the IEEE MIPR 2022.



SIWEI LYU (Fellow, IEEE) received the B.S. degree in information science and the M.S. degree in computer science from Peking University, China, in 1997 and 2000, respectively, and the Ph.D. degree in computer science from the Dartmouth College, in 2005. He is currently an SUNY Empire Innovation Professor at the Department of Computer Science and Engineering, the Director of the UB Media Forensic Laboratory (UB MDFL), and the Founding Co-Director of the

Center for Information Integrity (CII), University at Buffalo, State University of New York. His research interests include digital media forensics, computer vision, and machine learning.