# Detect and Locate: Exposing Face Manipulation by Semantic- and Noise-level Telltales

Chenqi Kong, Baoliang Chen, Haoliang Li, *Member, IEEE*, Shiqi Wang, *Senior Member, IEEE*, Anderson Rocha, *Senior Member, IEEE*, and Sam Kwong, *Fellow, IEEE*

*Abstract*—The technological advancements of deep learning have enabled sophisticated face manipulation schemes, raising severe trust issues and security concerns in modern society. Generally speaking, detecting manipulated faces and locating the potentially altered regions are challenging tasks. Herein, we propose a conceptually simple but effective method to efficiently detect forged faces in an image while simultaneously locating the manipulated regions. The proposed scheme relies on a segmentation map that delivers meaningful high-level semantic information clues about the image. Furthermore, a noise map is estimated, playing a complementary role in capturing low-level clues and subsequently empowering decision-making. Finally, the features from these two modules are combined to distinguish fake faces. Extensive experiments show that the proposed model achieves state-of-the-art detection accuracy and remarkable localization performance.

*Index Terms*—Face forensics, face forgery detection, face manipulation localization.

## I. INTRODUCTION

SUCCESSFUL computer vision technologies have recently also fueled face manipulation strategies with a series of methods proposed, including Deepfakes [1], Face2Face [2], FaceSwap [3] and NeuralTextures [4]. The fake content created by these methods is becoming increasingly realistic. Falsified video content, involving both manipulated identity and edited expression, raises various disconcerting problems within widespread social media, such as identity theft, fake news dissemination, and fraud. Even worse, most existing face manipulation technologies are disclosed online and easy to implement. Even a non-expert person without any prior professional skills can create high-quality fake content. It cannot be denied that these manipulation techniques and off-the-shelf commercial products (e.g., FakeApp [5]) bring certain positive impacts to the film-making industry and other entertainment applications. However, face manipulation techniques are highly likely to be misused, and the potential malicious applications are hazardous. As such, there has been an exponential increase in the demand for efficient face manipulation detection methods to counteract its dangerous impacts.

C. Kong, B. Chen, S. Wang and S. Kwong are with the Department of Computer Science, City University of Hong Kong, Hong Kong, China. (email: cqkong2-c@my.cityu.edu.hk; blchen6-c@my.cityu.edu.hk; shiqwang@cityu.edu.hk; cssamk@cityu.edu.hk).

H. Li is with the Department of Electrical Engineering, City University of Hong Kong, Hong Kong, China. (email: haoliang.li@cityu.edu.hk).

A. Rocha is with the Artificial Intelligence Lab. (Recod.ai) at the University of Campinas, Campinas 13084-851, Brazil (e-mail: arrocha@unicamp.br), URL: http://recod.ai
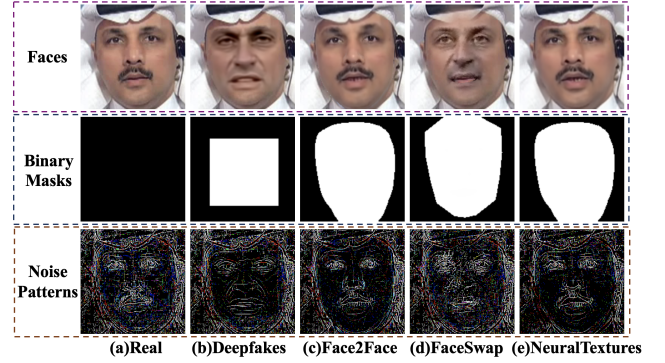
S. Wang is the corresponding author.

Fig. 1. Illustration of an authentic and four corresponding manipulated faces, including the face images (top row), binary masks (middle row), and noise patterns (bottom row). In the binary mask, the black pixels denote the masking label of the pristine face, and the white pixels represent the region that has been tampered with.

Despite the demonstrated success in developing detection methods, the advancement of face manipulation techniques still poses a grand challenge. The situation worsens when manipulated videos are compressed before distribution. One general detection methodology includes leveraging high-level semantic clues (e.g., lack of eye blinking [6], head pose inconsistency [7], and face splicing [8]). Such methods may ignore low-level signal variations in manipulated videos, resulting in limitations in terms of robustness and effectiveness. This has motivated the development of frequency-domain features [9]–[11]. Manipulation region localization [8], [12], [13] is another task of paramount importance in face forensics as it is pivotal to unveil the intention of a forger. However, this aspect has been largely ignored in most existing methods. In this paper, we conduct the manipulation localization at the pixel level. To produce manipulated faces, most face manipulation methods need to blend the altered face region into a background image. As such, the pixels in the manipulated face region will be altered due to the synthesis algorithm and the blending process. Another concern of face manipulation is the degradation of low-quality videos' detection accuracy due to the limited generalization capability across different quality levels. This could primarily affect the deployment of detection methods in practical applications.

In this vein, we aim to develop an method to expose face manipulation with three desired properties,

- **High detection accuracy**: It must achieve high accuracy in detecting manipulated faces.
- **Outstanding localization ability**: It must be able to

locate the manipulated regions precisely.

- **Flexibility**: It must be robust to videos in different quality levels.

We introduce a two-stream multi-scale framework based on the fusion of semantic-level and noise-level features, achieving the three desired properties mentioned above, simultaneously.

As Fig. 1 depicts, the semantic-level guidance (middle row) reveals the location of the manipulated face, rooted at the generally adopted procedure in face manipulation by blending the new face with a background. On the other hand, the noise-level guidance originates from noise patterns (bottom row), and observers may find that the noise pattern of the pristine face is more consistent and homogeneous than others. As such, semantic masks and noise patterns are employed as guiding labels to supervise the model's training. Besides, we extract three feature maps from the shallow, middle and deep layers of the convolutional neural network backbone, such that the extracted multi-scale features can carry information from low, middle, and high levels. Finally, we combine the semantic-level and noise-level features to check authenticity and locate the manipulated region of the input face. Extensive experiments demonstrate that the proposed framework achieves consistent performance improvement compared to state-of-the-art methods, holds promise for both high and low-quality videos, and delivers more information regarding the manipulated regions.

## II. RELATED WORK

### A. Face Manipulation and Detection

Generally speaking, face manipulation methods can be classified into the following two categories: identity swap and facial expression manipulation. In particular, Deepfakes [1] and FaceSwap [2] are two prominent representatives of face identity swap techniques. Face2Face [3] and NeuralTextures [4] are two popular facial expression editing methods. Deepfakes and NeuralTextures are learning-based methods among those four face manipulation techniques, while FaceSwap and Face2Face are computer graphics-based methods.

Due to the potential malicious applications of face manipulation techniques, numerous face forgery detectors have been proposed. For example, Li *et al.* [8] proposed the Face X-ray to detect the blending boundary of the fake face, which also reveals the manipulated boundaries. Chintha *et al.* [14] applied a recurrent structure to capture spatial and temporal information of manipulated videos. In [15] and [12], the authors designed two models to perform face forensics detection and localization. Li *et al.* [16] proposed a sharp-MIL (Multiple Instance Learning) to perform face forgery detection on videos. Qian *et al.* [9] introduced frequency analysis into face forgery detection and achieved promising detection results. Masi *et al.* [11] designed a two-branch recurrent network to detect face manipulation. None of the above methods considered combining semantic and noise signatures to detect forgeries and basically all of them suffer from performance drop when dealing with low-quality videos or cross-dataset evaluations.

### B. Binary Mask Supervision

Image manipulation segmentation focuses on regions which are potentially manipulated [17] and often relies upon guiding binary masks. The ground-truth masks are generated by marking forged regions as '1' and the pristine (un-forged) regions as '0'. In image forensics, Zhou *et al.* [18] fused the features from the RGB stream and the noise stream to capture the manipulated boundaries as well as the noise inconsistencies between the authentic and manipulated regions, and the features are then used to locate the potentially manipulated regions. Bappy *et al.* [17] designed a hybrid CNN-LSTM model to capture boundary discrepancy features and perform image manipulation localization. Zhou *et al.* [19] employed a novel generator to perform data augmentation and trained a model to complete image forgery localization. The binary mask guidance has also been widely used in face anti-spoofing. In [20]–[22], binary masks are employed to perform pixel-wise supervision during the training process. These models aim to discover arbitrary cues that can identify the live (pristine) and spoofed (manipulated) faces. In face manipulation detection, Du *et al.* [13] designed a locality-aware autoencoder to enforce the model to learn the feature representation of the forgery region, which boosts the generalization capability. Dang *et al.* [12] utilized the binary mask and an attention mechanism to make the extracted feature maps focus on manipulated regions, which further improves the detection performance. Li *et al.* [8] designed a more general model to capture the boundary splicing artifacts which can explicitly locate the potentially manipulated facial region.

Differently from the previous methods, in this paper, we design a multi-scale semantic map prediction module to capture the semantic-level information of input faces and simultaneously perform the localization of face manipulation. The binary masks play a supervision role in the model training process, encouraging the network to learn features that account for such fine-grained manipulations.

### C. Noise Modeling

Pioneering work on noise modeling for image forensics date back to 2004, where inconsistencies at the noise level have been leveraged to expose splicing artifacts of manipulated images [23]–[25]. More recently, photo-response non-uniformity (PRNU) [26] noise patterns have been widely studied in multimedia forensics. In general, PRNU can also be regarded as the device fingerprint. Korus *et al.* [27] proposed a multi-scale PRNU-based scheme to tackle the tampering localization problem. Quan *et al.* [28] focused on the correlation between the noise residual with the PRNU pattern for forgery detection. Cozzolino *et al.* [29] applied a noise extraction Siamese network [30] to achieve manipulation detection and forgery localization.

Noise patterns are also regarded as key clues in face anti-spoofing. More specifically, Jourabloo *et al.* [31] designed a CNN to decompose a spoofing face into the spoofing noise and a live face, and the spoofing noise is further used to perform the live/spoofing identification. Ren *et al.* [32] designed a
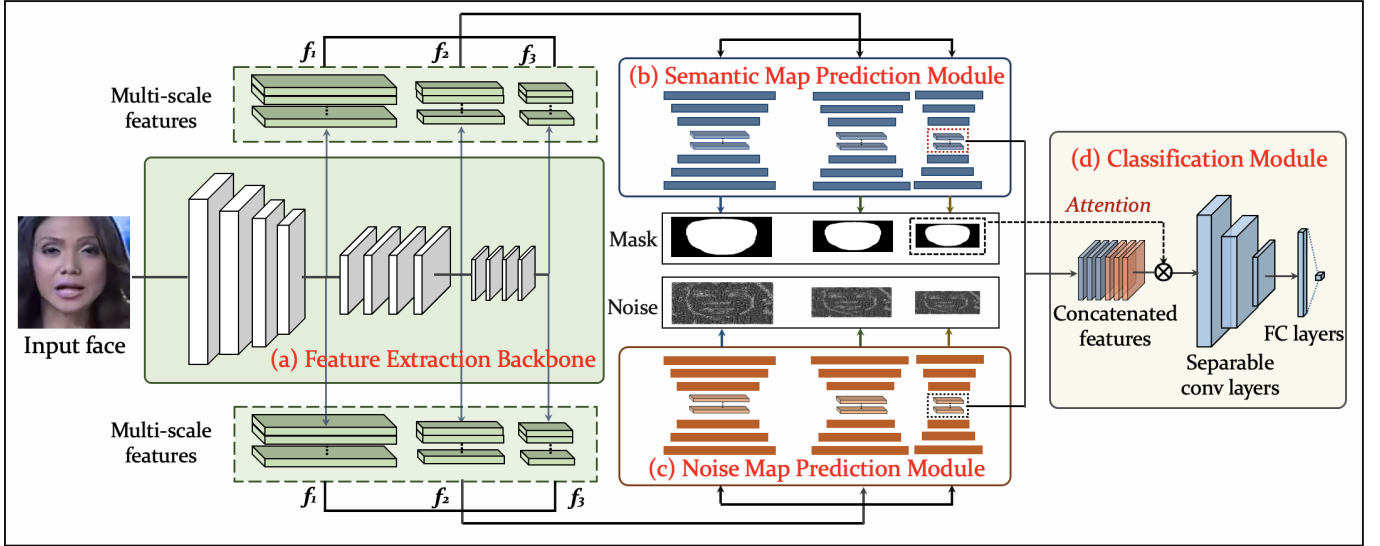
Fig. 2. Overview of the proposed framework. The face manipulation detection model consists of four components: (a) A feature extraction backbone; (b) A semantic map prediction module; (c) A noise map prediction module; and (d) a classification module. The faces are first forwarded to the feature extraction network backbone, and three feature maps extracted from the shallow, middle, and deep layers are synchronously processed by the semantic map prediction module and the noise map prediction module (the noise map intensity has been enlarged by five times for better visualization). A weighted summation over the three predicted semantic segmentation maps are conducted to perform the face manipulation localization. The estimated semantic and noise features are combined, and the spatial attention is conducted by element-wisely multiplying the combined features with the last predicted semantic segmentation maps. Finally, the attention-based feature maps are forwarded to the classification module for final decision-making.

noise-attention network architecture to extract noise-features for live/spoofing face classification.

While noise signature has been widely used in image forensics and face anti-spoofing tasks, how the noise signature benefits face manipulation detection, especially the fake face contents generated by deep learning technologies, has not been investigated yet. In this work, we argue that some key cues remain in the noise maps of manipulated faces. Therefore, a noise map prediction module is designed to extract the noise pattern of corresponding input faces, and the noise feature is subsequently fed to a classification module for face manipulation detection.

## III. PROPOSED METHOD

This section presents a method to simultaneously detect and locate face manipulated regions. It relies upon high-level semantic information clues about an image combined with noise low-level features. The method consists of four components: (a) a feature extraction backbone; (b) a semantic map prediction module; (c) a noise map prediction module; and (d) a classification module.

### A. Overall Framework

We propose a two-stream multi-scale framework for face forgery detection and localization, and the architecture is illustrated in Fig. 2. In particular, we denote the training dataset as $D = \{F_i, M_i, n_i, c_i\}_{i=1}^N$, where $N$ is the total number of training samples. The dataset consists of the following four components: the $i$th input face $F_i$, its corresponding binary mask label $M_i$, noise pattern label $n_i$ and ground-truth binary label of forgery $c_i$. We employ the Xception network [33] as

the backbone to learn multi-scale feature maps based on two insights. Other networks could also be used without loss of generality.

First, multi-scale features enable the model to learn both semantic and geometric information as features from different layers contribute to different receptive fields. Second, multi-scale supervision leads the model to focus on the valuable information from the beginning (i.e., learning face-forgery-specific information at the shallow layer), which benefits face manipulation detection and localization performance.

The semantic map prediction module and the noise map prediction module synchronously processes the multi-scale feature maps, where the outputs are guided by binary mask label $M_i$ and noise pattern label $n_i$, respectively. The semantic map prediction module is trained to capture the high-level semantic information from the extracted multi-scale features where the estimated semantic map indicates the potentially manipulated regions in the corresponding input faces.

The three predicted semantic segmentation maps are leveraged to perform forgery localization. The noise map prediction module is further proposed to enforce the multi-scale features to focus on image content and pay attention to content-irrelevant low-level information. The estimated noise maps contain rich high-frequency cues that might expose noise artifacts in manipulated faces by noise modeling.

We combine the last semantic features and the last noise features, and the spatial attention is sequentially conducted by element-wisely multiplying the concatenated features with the last predicted semantic segmentation maps. Finally, the attention-based features are forwarded to the classification module to identify the authenticity of input faces.

The overall objective function consists of three components:

$$L = L_c + \lambda_1 L_n + \lambda_2 L_b, \qquad (1)$$

where $L_c$, $L_n$, and $L_b$ denote final classification loss, noise map prediction loss, and semantic map estimation loss, respectively. $\lambda_1$ and $\lambda_2$ are hyper-parameters to weigh the loss components. $L_c$ is the cross-entropy loss between the prediction result $\hat{c}_i$ and ground truth label $c_i$:

$$L_c = -\frac{1}{N} \sum_{i=1}^{N} (c_i \log \hat{c}_i + (1 - c_i) \log(1 - \hat{c}_i)), \qquad (2)$$

We better describe $L_n$ and $L_b$ in the following subsections.

### B. Semantic Map Prediction Module

Binary mask supervision has been widely used to tackle various forensics problems. In this paper, we generate the ground-truth mask for each face in the training stage and use the mask to supervise the training of the proposed model. The benefits of the semantic map prediction module are mainly threefold: 1) it enables the model to localize manipulated regions, providing evidence to show whether the input faces have been manipulated; 2) it constrains the model to focus on manipulated regions, leading the model to achieve a better manipulation detection performance; and 3) different manipulation approaches tend to have different binary mask shapes, thus providing auxiliary information to complete multi-class classification (more details are elaborated in Sec.IV-C.4)).

The semantic map prediction module takes the extracted multi-scale features as inputs, where the outputs are the estimated semantic segmentation maps. In particular, we apply the depth-wise separable convolution block as our semantic map prediction module as it has been shown to be effective in different types of problems [33]. For the output of the semantic map prediction module, each pixel value indicates the probability that the corresponding receptive field is a manipulated region in the input face. Herein, we conduct a weighted summation strategy over the estimated semantic maps in multi-scale to determine the final manipulation localization map of the input faces.

We use videos in FaceForensics++ [34] as our training set as it provides a manipulation mask for each manipulated face, supervising the training of the semantic map prediction module. Training the semantic map prediction module can tackle a binary classification problem, including the manipulated region and the non-manipulated area. Furthermore, the cross-entropy loss is leveraged to supervise the estimation of semantic segmentation maps:

$$L_b = -\frac{1}{N} \sum_{i=1}^{N} \sum_{j} \frac{1}{N_p} \sum_{x,y} (M_{i,j}^{x,y} \log \hat{M}_{i,j}^{x,y} + (1 - M_{i,j}^{x,y}) \log(1 - \hat{M}_{i,j}^{x,y})), \qquad (3)$$

where $\hat{M}_{i,j}^{x,y}$ and $M_{i,j}^{x,y}$ represent the estimated and ground truth semantic maps, and the sizes of all ground truth maps have been aligned to the corresponding predicted maps. The parameters $x$ and $y$ denote the pixel location; $i$ specifies $i$th
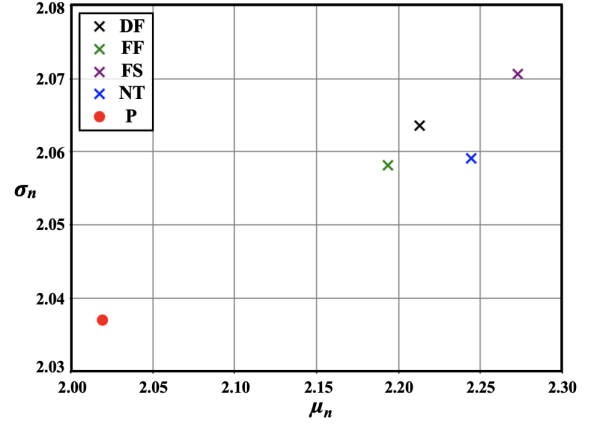


Fig. 3. The statistical analysis of faces from four manipulation methods and pristine faces. Red color dots represent pristine (P) face images, and the remaining four colors indicate the four manipulation methods: Deepfakes(DF), Face2Face(FF), FaceSwap(FS), and NeuralTextures(NT).

input face; $j = \{shallow, middle, deep\}$ denotes the mask predicted by the specific feature layer; $N_p$ is the total pixel number.

### C. Noise Map Prediction Module

One notable example of forensics-related noise pattern analysis is PRNU, which is the dominant part of the noise in natural images caused by the inhomogeneity of silicon wafers and imperfections during the sensor manufacturing process [35] and has been widely applied in different multimedia forensics tasks [27], [29], [35] to date. Inspired by the PRNU paradigm, we leverage the wavelet-based filter [36] to extract a noise map to explore further the face image forensics artifacts as the low-level clue guidance for face manipulation detection. The noise map prediction module is designed to make the extracted multi-scale features irrelevant with high-level semantic content (i.e., facial appearance) and provide a complementary clue for face manipulation detection.

The design of the noise map prediction module relies upon the assumption that noise patterns of manipulated faces expose two artifacts. First, the noise pattern of the manipulated region tends to expose more abnormal artifacts due to the image processing of the fake face creation process. Second, the statistical discrepancy between the noise distributions of the manipulated face region and the background region introduces noise distribution inconsistencies [18], [29]. We statistically analyze the noise pattern distributions of four manipulation methods and pristine data in Fig. 3. It can be observed that the mean $\mu_n$ and variance $\sigma_n$ values of pristine faces are much smaller than the values of the other four manipulation methods. Therefore, it is reasonable to employ noise guidance to perform face manipulation detection.

Following previous multimedia forensics methods [27], [29], [35], we apply a wavelet-based filter [36] to extract noise patterns for each face, which can be formulated as:
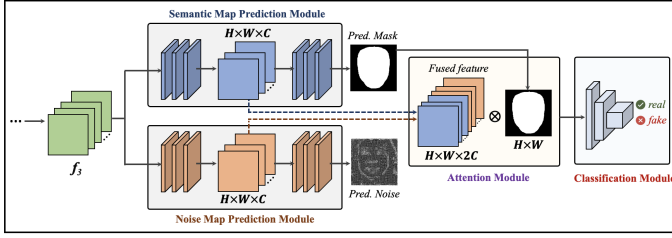
$$n_i = F_i - f(F_i) \qquad (4)$$

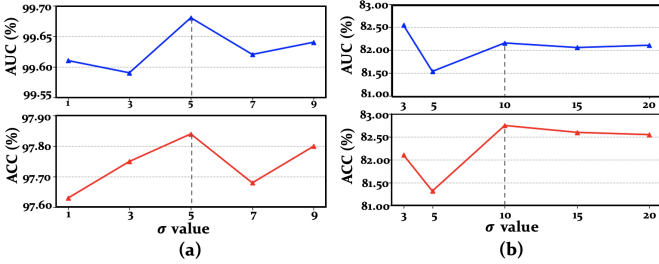Fig. 4. Illustration of feature fusion strategy and attention mechanism.



Fig. 5. The manipulation detection performances of noise prediction module on (a) high-quality faces and (b) low-quality faces. The ground-truth noise maps are generated with different $\sigma$ values.

where $F_i$ and $n_i$ represent the $i$th input raw face and the corresponding extracted noise pattern, and $f(\cdot)$ is defined as the noise filter:

$$\hat{F}_i = f(F_i) = \frac{\hat{\sigma}_{F_i}^2}{\hat{\sigma}_{F_i}^2 + \sigma^2} F_i \qquad (5)$$

where $\hat{F}_i$ is the denoised face image, $\sigma^2$ indicates the variance estimation of the additive white Gaussian noise (AWGN), and $\hat{\sigma}_{F_i}^2$ is the variance of $\hat{F}_i$. The estimation details of $\hat{\sigma}_{F_i}^2$ can be found in [36].

We can extract the noise residual for each input face as the noise label. We further design a multi-scale noise map prediction module to estimate low-level noise patterns and extract noise-specific features. The extracted noise-specific features aim to expose noise artifacts, serving as auxiliary information to improve the final authentic identification performance. The noise map prediction loss $L_n$ calculates the $l_1$ norm of the pixel-wise difference between predicted noise maps $\hat{n}$ and corresponding noise labels $n$:

$$L_n = \frac{1}{N} \sum_{i=1}^{N} \sum_{j} \|n_{i,j} - \hat{n}_{i,j}\|_1, \qquad (6)$$

where $i$ specifies $i$th input face; $j$ denotes the specific feature layer; $n_{i,j}$ and $\hat{n}_{i,j}$ represent the ground truth and estimated noise patterns. This design enables the network to learn the content-irrelevant low-level information and expose noise artifacts.

### D. Feature Fusion and Classification Module

In this subsection, we present the feature fusion process, attention mechanism, and the final classification module. We illustrate the feature fusion strategy and the attention mechanism in Fig. 4. $f_3$ represents the feature extracted from the deep layer of the backbone. First, we feed $f_3$ forward to Semantic Map Prediction Module and Noise Map Prediction Module to obtain the predicted mask and noise map. The high-level semantic feature and low-level noise feature have the same size of H×W×C, and the predicted map size is H×W. Next, we combine the semantic and noise features by concatenating semantic features and low-level features along the channel dimension. Subsequently, we conduct the spatial attention by element-wisely multiplying the fused feature with the predicted mask. Intuitively, the manipulated face region contains abundant abnormalities and artifacts. We aim to mine the artifacts at both semantic- and noise-levels. Therefore, we design the attention mechanism to highlight the manipulated locations and further empower the final binary classification performance. On the other hand, each pixel in the attention map indicates the probability that the corresponding receptive field is a manipulated region in the input face. Element-wisely multiplying the combined features with the attention map can guide the model to focus on the manipulated spatial region and improve binary classification accuracy. Finally, the spatially attended feature is fed to the Classification Module to identify the authenticity of the input query.

## IV. EXPERIMENTS

### A. Data Preparation

*1) Dataset:* In this paper, we conduct the experiments on the challenging FF++ [34] dataset. FF++ [34] is a dataset of face manipulation that contains 1,000 pristine videos and 4,000 associated manipulated videos created by four state-of-the-art forgery techniques: Deepfakes [1], Face2Face [2], FaceSwap [3] and NeuralTextures [4]. Besides, FF++ [34] provides three quality levels controlled by the quantization parameters (QP) in compression for these 5000 videos: raw (QP=0), HQ (high-quality, QP=23), and LQ (low-quality, QP=40). Considering the deployment in real-world application scenarios, we conduct our experiments on both HQ videos and LQ videos. FF++ [34] also provides the ground truth-masks that indicate the forged regions of manipulated faces, enabling the developments of face forgery localization methods. Following the experimental setting in [34], we take 720 videos for training, 140 videos for validation, and 140 videos for testing. We extract 270 frames from each training video and 100 frames from each validation and testing video. Following the official experimental setting [34], during training we augment the number of real faces four times to address the data imbalance between real and fake faces. The summary of the data in training, validation, and testing is listed in Table I. For that we just extract more faces from the training videos.

*2) Binary mask generation:* FF++ provides a 3D mask for every video frame, and the 3D mask indicates the manipulated region of each face. As such, we generate the ground-truth binary mask for each face according to the boundary location of the 3D mask. The middle row in Fig. 1 illustrates the generated ground-truth masks for real and manipulated faces, where the white pixels denote the manipulated regions, and the black pixels represent the non-manipulated regions (pristine regions). For other data corpora which do not provide mask

TABLE I
DATASET BREAKDOWN FOR TRAINING, VALIDATION, AND TESTING SETS.

| Methods | Training | | Validation | | Test | |
|---|---|---|---|---|---|---|
| | Videos | Faces | Videos | Faces | Videos | Faces |
| Deepfakes | 720 | 192,748 | 140 | 13,907 | 140 | 14,000 |
| Face2Face | 720 | 194,091 | 140 | 14,000 | 140 | 14,000 |
| FaceSwap | 720 | 193,989 | 140 | 14,000 | 140 | 14,000 |
| NeuralTextures | 720 | 194,004 | 140 | 14,000 | 140 | 14,000 |
| Real | 720 | 776,912 | 140 | 56,000 | 140 | 56,000 |

labels, the alternative way to quickly generate manipulation masks manually can be summarized as two steps: (1) computing the absolute difference between the fake face and the corresponding real face, as most datasets contain the real and fake video pairs, *e.g.*, CelebDF [37], DFD [38], and DFDC [39]; (2) setting a threshold value for the absolute difference map and convert it to a binary mask. This mask generation process can be summarized as follows,

$$M_{i,j} = \begin{cases} 0, & |F_{i,j}^{real} - F_{i,j}^{fake}| < threshold, \\ 1, & |F_{i,j}^{real} - F_{i,j}^{fake}| \geq threshold. \end{cases} \quad (7)$$

*3) Noise map generation:* Following previous forensics methods [27], [29], [35], we apply the wavelet-based noise filter [36] to extract the noise map for each input face. The extracted noise maps are regarded as pseudo labels to supervise the training of the noise map prediction module. The hyperparameter $\sigma$ as defined in Eq. (5) is the variance estimation of additive white Gaussian noise(AWGN). To determine the best values of $\sigma$ for high-quality and low-quality faces, we leverage the noise patterns generated with different $\sigma$ values to supervise the training of the noise map prediction module and the detection performances of high-quality and low-quality faces are illustrated in Fig. 5. According to the detection results in Fig. 5 (a) and (b), we set the $\sigma$ values as 5 and 10 for high-quality and low-quality faces.

### B. Implementation Details

*1) Training strategy:* We apply the popular `dlib` face detector [40] to crop the face regions enlarged by a factor of 1.3. The proposed framework is implemented by Pytorch [41]. The model is trained using Adam optimizer [42] with $\beta_1$=0.9 and $\beta_2$=0.999. Multi-scale feature layers are extracted after the conv2 , block2, and block11 of the Xception backbone. We release the code here[1].

We conduct a two-step training strategy: 1) training the Xception backbone and the semantic map prediction module, where the weights of the backbone are initialized with the ImageNet [43] weights; and 2) initializing the weights of the Xception backbone and the semantic map prediction module with the weights obtained in Step 1), then training the overall framework in an end-to-end manner. This two-step training strategy can effectively alleviate the over-fitting problem and be more robust to local minima. We set the learning rate and weight decay as 0.0002 and 1e-5, respectively. The model is trained on 2 RTX 2080Ti GPUs with batch size 32.

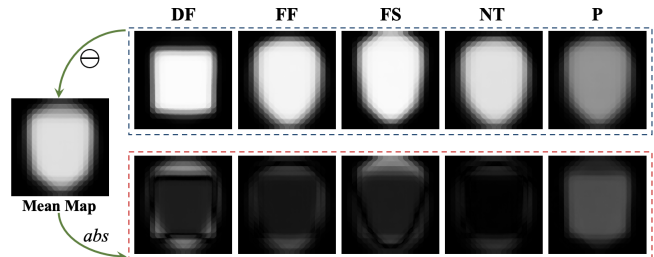[1]https://github.com/ChenqiKONG/Detect_and_Locate.



Fig. 6. Predicted mask maps from the testing set. The mean map on the left represents the average summation over all manipulation localization maps and pristine maps. The top row shows the mean maps of four specific manipulation methods (DF: Deepfakes, FF: Face2Face, FS: FaceSwap, NT: NeuralTextures) and pristine face images (P: Pristine). The maps on the bottom row indicate the absolute differences between the top-row maps and the left mean map.
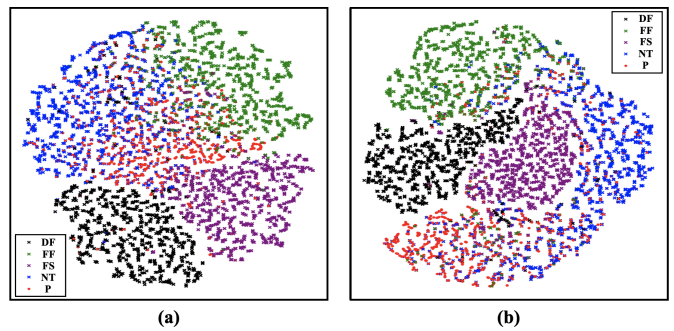


Fig. 7. T-SNE feature embedding visualization of (a) the Xception baseline and (b) the proposed method on FF++ [34] low-quality (QP=40) dataset in the five-class classification task. Red color dots represent pristine (P) face images, and the rest four colors indicate the four manipulation methods: Deepfakes (DF), Face2Face (FF), FaceSwap (FS), and NeuralTextures (NT).

*2) Evaluation metrics:* Following most existing face forgery detection methods, we employ the detection Accuracy rate (ACC) and Area Under the Receiver Operating Characteristic Curve (AUC) as the forgery detection evaluation metrics. ACC is the most straightforward metric in face forgery detection, which directly reflects the detection capability of the detector. AUC is a more objective evaluation metric, which has been widely applied in recent state-of-the-art detection methods, such as Face X-ray [8], Celeb-DF [37], and Two-branch [11].

This paper takes AUC as our key evaluation metric and reports all detection results at the frame level. We further use the metrics of False Positive Rate (FPR), False Negative Rate (FNR), Equal Error Rate (EER), and Average Precision (AP) to interpret the face manipulation detection results better.

On the other hand, we utilize Intersection over Union (IoU), Pixel-wise Binary Classification Accuracy (PBCA), and Inverse Intersection Non-Containment (IINC) [12] as the localization evaluation metrics. IoU is one of the standard metrics in semantic segmentation and has been widely used in numerous semantic segmentation tasks [44], [45]. The definition of IoU is illustrated as follows,

$$IoU = \frac{\hat{M} \cap M}{\hat{M} \cup M}. \quad (8)$$

The numerator calculates the area of overlap between the

predicted mask $\hat{M}$ and the ground truth mask $M$, and the denominator represents the area encompassed by both the predicted mask and the ground truth mask. In addition to IoU, PBCA and IINC proposed in [12], are two metrics for evaluating face manipulation localization performance. PBCA calculates the mask prediction accuracy at the pixel level. IINC measures the non-overlap ratio of both predicted map and ground truth map. Thus it is a more robust and coherent metric in evaluating the localization results. As such, IoU, PBCA, and IINC are employed to evaluate the forgery localization accuracy.

*C. Manipulation Detection*

We compare the proposed method with previous detection methods in terms of ACC, EER, and AUC.

*1) Detection performance on all manipulation methods:* Table II reports the detection results of the proposed framework and previous methods trained on the FF++ dataset. We use the bold font to highlight the best results while underlining the second-best results among all listed methods. HQ and LQ indicate the high-quality and low-quality data. Comparing with the most recent work [46], the proposed model achieves a significant ACC score improvement, going from 80.18% → 84.84%. Moreover, the overall detection performance of the proposed method outperforms all listed methods for HQ and LQ image qualities. These performance gains on both ACC and AUC scores are mainly due the the high-level and low-level joint learning strategy adopted in our formulation, further demonstrating the effectiveness of the proposed model.

*2) Detection accuracy on specific manipulation methods:* Detecting low-quality manipulated face images is a challenging task as severe compression erases much detailed information from the original faces. To demonstrate that the proposed model can achieve remarkable detection performance on low-quality faces, we list the detection accuracy (ACC) of specific manipulation methods on the low-quality FF++ dataset in Table III. DF, FF, FS, and NT represent four manipulation methods: Deepfakes, Face2Face, FaceSwap, and NeuralTextures. The manipulation-specific detectors are trained and tested on the same face manipulation methods. Comparing with previous detection methods, the proposed model achieves the best detection accuracy on all four manipulation methods.

*3) Cross-dataset evaluation:* Most existing detection models always suffer a significant performance drop when applied to unseen datasets. To comprehensively evaluate the generalization ability of the proposed model, we conduct extensive cross-dataset evaluation experiments in this paper. We train our model on the Deepfakes (HQ and LQ) and Pristine (HQ and LQ) data of FF++ and test it on the unseen CelebDF [37], DFD [38], DFDC [39], and Deepfake-TIMIT [53] datasets. For a fair comparison, we fully align the training, validation, and testing data. We list the dataset split strategy in Table IV. The training set includes 720 HQ, 720 LQ Deepfake videos, and 720 HQ, 720 LQ Pristine videos in FF++, and 270 frames are extracted from each training video. For the validation and test set (intra), we extract 100 frames for each video. For the unseen Deepfake datasets, we extract 30 frames from each video at equal intervals.

TABLE II
DETECTION PERFORMANCE TRAINED AND EVALUATED ON ALL MANIPULATED AND PRISTINE DATA. HQ AND LQ DENOTE THE HIGH-QUALITY AND LOW-QUALITY DATA. COMPARED TO PREVIOUS DETECTION METHODS, THE PROPOSED APPROACH ACHIEVES SUPERIOR QUANTITATIVE RESULTS FOR HIGH-QUALITY AND LOW-QUALITY FACE IMAGES.

| Methods | LQ (QP=40) | | HQ (QP=23) | |
|---|---|---|---|---|
| | AUC(%) | ACC(%) | AUC(%) | ACC(%) |
| Steg. Features [47] | - | 55.98 | - | 70.97 |
| Cozzolino *et al.* [48] | - | 58.69 | - | 78.45 |
| Bayar & Stamm [49] | - | 66.84 | - | 82.97 |
| Rahmouni *et al.* [50] | - | 61.18 | - | 79.08 |
| MesoNet [51] | - | 70.47 | - | 83.10 |
| Xception [33] | 84.38 | 83.61 | 99.71 | 98.04 |
| Face X-ray [8] | 61.60 | - | 87.35 | - |
| Two-branch [11] | 86.59 | - | 98.70 | - |
| SPSL [10] | 82.82 | 81.57 | 95.32 | 91.50 |
| Nirkin *et al.* [46] | - | 80.18 | - | - |
| F$^3$Net [9] | **87.26** | <u>84.65</u> | 99.60 | 97.49 |
| Multi-Att [52] | 85.52 | 84.49 | <u>99.75</u> | <u>98.33</u> |
| Ours | <u>87.10</u> | **84.84** | **99.77** | **98.40** |

The intra- and cross-dataset evaluation results (AUC & EER) are listed in Table V. For all detection methods listed in Table V, we select the checkpoint with the best AUC score on the validation set and test it on all test sets. In this part, we apply the designed two-stream learning scheme on both Xception and EfficientNet backbones, which are respectively denoted as Ours (Xcep.) and Ours (Effi.) in Table V. Our frameworks achieve the best AUC and EER on the FF++/DF (HQ) set and the FF++/DF (LQ) set. Under the cross-dataset evaluation setting, the proposed method outperforms all the listed models in terms of the average AUC and EER values. The cross-dataset evaluation experiment demonstrates that the proposed model is capable of achieving high generalization capability. Compared with the Xception baseline, our framework gains a 6.72% AUC and 4.76% EER average score improvement under the cross evaluation settings. This indicates that the proposed two-stream multi-scale learning framework can also boost the backbone model's generalization ability.

*4) Multi-class classification evaluation:* Although identifying the authenticity of input faces is of great importance, specifying the manipulation method is also a non-trivial problem. However, correctly classifying the manipulation approach of the input face can reveal whether the identity or expression of the input face has been manipulated, which further unveils the potential intent of the forger.

We further evaluate the proposed model on this five-way (real and four respective manipulation methods) classification task. The classification results on the FF++ low quality (QP=40) dataset are reported in Table VI. We use the bold font to highlight the best results while underlining the second-best results among all listed methods.

Compared with the result of the Xception baseline, the proposed model equipped with the semantic segmentation module (Xception(seg)) in the second last row achieves a 5.61% recall rate improvement, going from 75.43% → 81.04%. We argue that the classification performance gain mainly benefits from
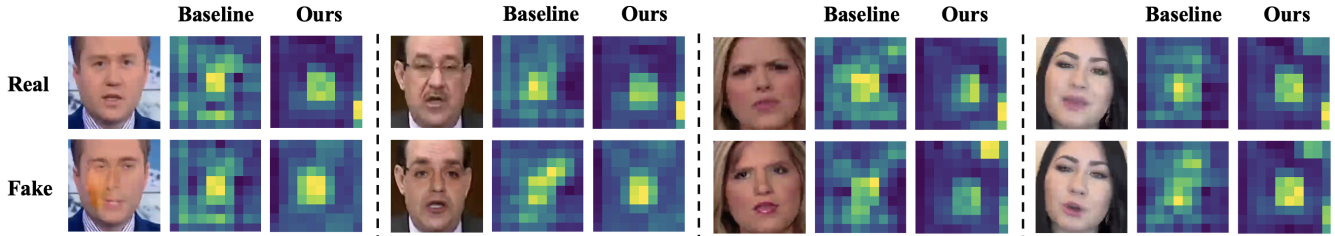
Fig. 8. Feature map visualization of the Xception baseline and the proposed method.

TABLE III
DETECTION ACCURACY (ACC) OF SPECIFIC MANIPULATION METHODS ON
LOW-QUALITY FACES (DF: DEEPFAKES, FF: FACE2FACE, FS: FACESWAP,
NT: NEURALTEXTURES).

| Methods | DF(%) | FF(%) | FS(%) | NT(%) |
|---|---|---|---|---|
| Steg. Features [47] | 67.00 | 48.00 | 49.00 | 56.00 |
| Cozzolino *et al.* [48] | 75.00 | 56.00 | 51.00 | 62.00 |
| Bayar & Stamm [49] | 87.00 | 82.00 | 74.00 | 74.00 |
| Rahmouni *et al.* [50] | 80.00 | 62.00 | 59.00 | 59.00 |
| MesoNet [51] | 90.00 | 83.00 | 83.00 | 75.00 |
| SPSL [10] | 93.48 | 86.02 | 92.26 | 76.78 |
| Xception [33] | 97.16 | 91.02 | 96.71 | 82.88 |
| Ours | **97.25** | **94.46** | **97.13** | **84.63** |

TABLE IV
DATASET BREAKDOWN FOR TRAINING, VALIDATION, AND TESTING SETS
UNDER THE CROSS-DATASET EVALUATION SETTING.

| | | # Video | | # Frame | |
|---|---|---|---|---|---|
| | Dataset | Real | Fake | Real | Fake |
| **Training** | FF++ | 1,440 | 1,440 | 388,691 | 386,323 |
| **Validation** | FF++ | 280 | 280 | 28,000 | 27,818 |
| **Test (Intra)** | FF++ | 280 | 280 | 28,000 | 27,957 |
| **Test (Cross)** | CelebDF | 590 | 5,639 | 17,700 | 169,170 |
| | DFD (HQ) | 363 | 3,068 | 10,890 | 92,040 |
| | DFD (LQ) | 363 | 3,068 | 10,890 | 92,040 |
| | DFDC | 19,154 | 100,000 | 574,620 | 3,000,000 |
| | DF-TIMIT (HQ) | 240 | 320 | 7,200 | 9,600 |
| | DF-TIMIT (LQ) | 240 | 320 | 7,200 | 9,600 |

the supervision of the binary mask.

The benefit of the semantic map prediction module is twofold. First, this constrains our model to focus on the manipulated region of the input face, leading the model to mine more manipulation-related information. Second, different manipulation methods tend to leave different binary mask shapes, thus providing auxiliary information to perform the five-class classification.

As such, the scheme of binary mask supervision can separate the faces manipulated by different approaches. To better clarify this point, we further present the visualization of predicted mask maps on the testing set in Fig. 6. The mean map on the left represents the average summation of overall predicted localization maps and pristine maps. The top row shows the mean maps of four specific manipulation methods and pristine face images. The absolute difference maps between the top-row maps and the mean map are shown in the bottom row. It can be readily observed that each bias map is discriminative from the others.

The final classification result in the last row of Table VI shows that the proposed model outperforms the previous state-of-the-art method, SPSL, by 3.25% average recall rate, going from 78.94% → 82.19%. The noise map prediction module also mines significant manipulation method clues. Thus, we can observe that the final result (Xception(fusion)) achieves a 1.15% average recall rate and a 6.22% pristine recall rate improvement comparing with the results of Xception(seg), demonstrating the effectiveness of the noise map prediction module from a complementary viewpoint.

Furthermore, we show the t-SNE feature embedding visualization of the Xception baseline and the proposed method in Fig. 7. As shown in Fig. 7 (a), the Xception model cannot

separate different manipulation methods and pristine face images, and the pristine faces are clustered with NeuralTextures and Face2Face fake faces in the feature space. However, the proposed model achieves a better multi-class embedding division performance in the t-SNE feature space, and the pristine data is less confused with other manipulated data, demonstrating the effectiveness of the proposed model from another point of view.

*5) Feature map visualization:* To better demonstrate the effectiveness of the proposed model, we further visualize the feature maps of the Xception baseline and the proposed model ones in Fig. 8. All real and fake faces are from the FF++ low-quality dataset, and the models are trained under the setting of Sec.IV-C.1).

As shown in Fig. 8, the baseline feature maps of real and corresponding fake faces are similar, resulting in a struggling manipulation detection performance. Conversely, the proposed model prefers to focus on the central regions for fake faces and the peripheral regions for pristine faces. As a result, comparing with the Xception baseline, the real and fake feature maps extracted from the proposed model are more discriminative, thus leading to a better manipulation detection performance.

*6) Discussions:* In this subsection, we quantitatively report the face forgery detection performances on various experimental settings and qualitatively visualize the feature maps and t-SNE feature embedding distributions. The presented results demonstrate the effectiveness and robustness of the proposed framework, which can capture the artifacts in both semantic-level and noise-level, thus achieving superior detection performance on all mentioned experimental settings.
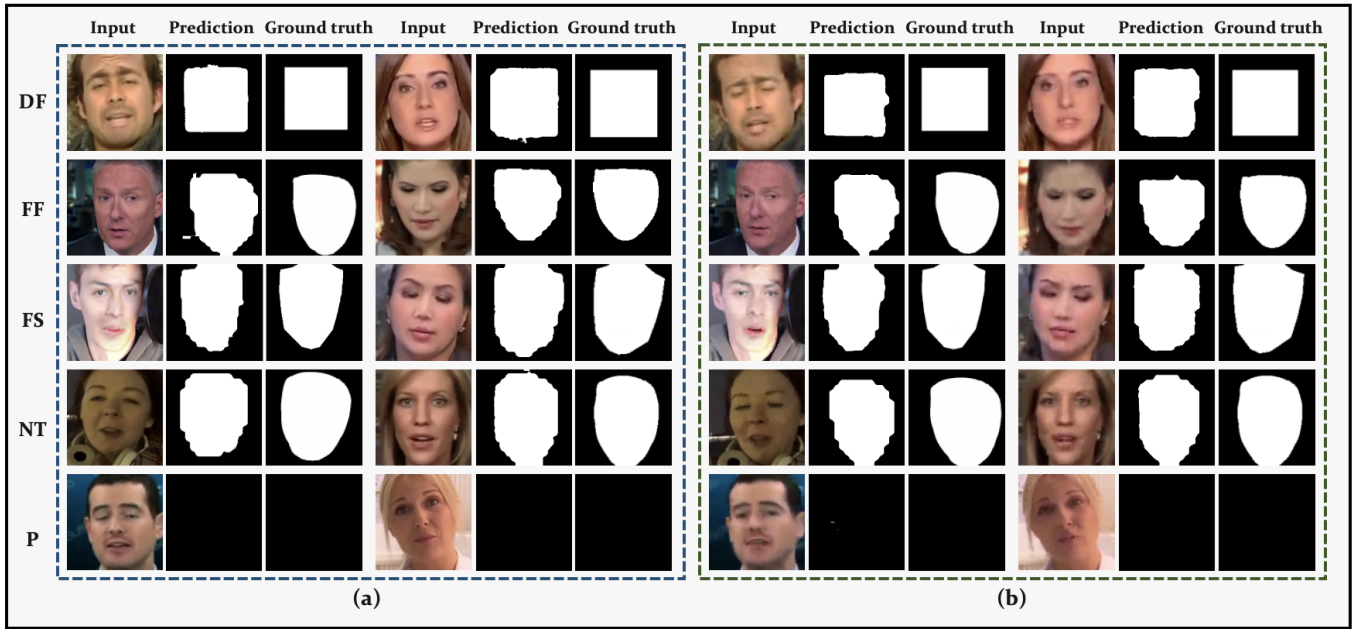
Fig. 9. Face manipulation localization results on (a) HQ faces (QP=23) and (b) LQ faces (QP=40). DF: Deepfakes, FF: Face2Face, FS: FaceSwap, NT: NeuralTextures, P: Pristine. The input face, predicted forgery localization map, and corresponding ground truth mask are presented in the left, middle, and right columns. It can be observed that the proposed method can locate the manipulated regions of the input faces with different head poses and various lighting conditions.

TABLE V
CROSS-DATASET EVALUATION RESULTS.

| Dataset | INTRA (%) | | | | CROSS (%) | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | FF++/DF (HQ) | | FF++/DF (LQ) | | CelebDF | | DFD (HQ) | | DFD (LQ) | | DFDC | | DF-TIMIT (HQ) | | DF-TIMIT (LQ) | | AVG (Cross) | |
| Method | AUC | EER | AUC | EER | AUC | EER | AUC | EER | AUC | EER | AUC | EER | AUC | EER | AUC | EER | AUC | EER |
| MesoNet [51] | 76.34 | 30.23 | 70.11 | 35.39 | 58.85 | 43.89 | 62.07 | 41.75 | 52.25 | 48.65 | 54.60 | 46.67 | 33.61 | 60.16 | 45.08 | 53.04 | 51.08 | 49.03 |
| MesoIncep4 [51] | 98.35 | 6.72 | 95.56 | 11.68 | 68.26 | 37.59 | 79.18 | 27.68 | 63.27 | 40.37 | 61.92 | 41.12 | 16.12 | 76.18 | 27.47 | 66.77 | 52.70 | 48.29 |
| ResNet50 [54] | 99.15 | 4.35 | 98.46 | 6.57 | 67.09 | 37.57 | 69.60 | 35.42 | 60.61 | 42.23 | 61.97 | 41.19 | 41.95 | 55.97 | 47.27 | 52.33 | 58.08 | 44.12 |
| Face X-ray [8] | 99.60 | 2.45 | 99.06 | 4.41 | 71.89 | 32.64 | 69.61 | 33.97 | 62.89 | 39.58 | 58.97 | 43.52 | 42.52 | 55.07 | 50.05 | 49.11 | 59.32 | 42.31 |
| DFFD [12] | 99.74 | 2.56 | 99.37 | 3.92 | 69.55 | 35.92 | 71.69 | 33.76 | 60.60 | 42.32 | 59.72 | 42.92 | 32.91 | 61.16 | 39.32 | 57.06 | 55.63 | 45.52 |
| Multi-task [55] | 98.57 | 5.94 | 96.02 | 10.62 | 65.18 | 38.92 | 70.75 | 34.55 | 58.61 | 44.49 | 57.38 | 44.83 | 16.53 | 77.86 | 15.59 | 78.50 | 47.34 | 53.19 |
| EfficientNet [56] | 99.65 | 1.93 | 99.19 | 4.06 | 75.90 | 30.94 | 80.63 | 27.07 | 64.19 | 39.87 | 66.39 | 38.68 | 29.12 | 65.75 | 28.34 | 57.43 | 57.43 | 44.59 |
| F$^3$Net [9] | 99.74 | 2.16 | 99.21 | 4.47 | 72.28 | 33.85 | 72.92 | 33.04 | 58.89 | 43.82 | 63.33 | 40.54 | 38.55 | 58.33 | 45.67 | 52.72 | 58.61 | 43.72 |
| Xception [33] | 99.64 | 2.78 | 99.16 | 4.23 | 67.75 | 37.26 | 72.45 | 33.50 | 59.73 | 43.12 | 63.12 | 40.58 | 33.82 | 62.83 | 40.79 | 57.44 | 56.28 | 45.79 |
| Ours (Effi.) | 99.73 | 1.67 | **99.46** | **2.92** | 67.15 | 37.93 | 73.52 | 33.06 | 67.21 | 37.37 | 60.32 | 43.12 | 49.90 | 49.47 | 51.01 | 48.03 | 61.52 | 41.50 |
| Ours (Xcep.) | **99.85** | **1.80** | 99.30 | 4.10 | 70.65 | 35.45 | 76.23 | 30.26 | 64.53 | 39.59 | 63.31 | 40.79 | 47.20 | 53.23 | 56.08 | 46.86 | **63.00** | **41.03** |

TABLE VI
RECALL RATES OF GENUINE AND FOUR MANIPULATION METHODS WITH C40(QP=40) SETTING IN THE FIVE-CLASS CLASSIFICATION. DF: DEEPFAKES, FF: FACE2FACE, FS: FACESWAP, NT: NEURALTEXTURES, P: PRISTINE. WE USE THE BOLD FONT TO HIGHLIGHT THE BEST RESULTS WHILE UNDERLINING THE SECOND-BEST RESULTS AMONG ALL LISTED METHODS.

| Methods | DF(%) | FF(%) | FS(%) | NT(%) | P(%) | AVG(%) |
|---|---|---|---|---|---|---|
| MesoNet [51] | 62.45 | 40.37 | 28.89 | 63.35 | 40.93 | 47.20 |
| ResNet50 [54] | 82.76 | 76.30 | 81.21 | 52.18 | 46.82 | 67.85 |
| Xception [33] | 86.61 | 78.88 | 83.16 | 52.94 | 75.55 | 75.43 |
| SPSL [10] | 91.16 | 78.31 | 88.75 | 58.97 | **77.49** | 78.94 |
| Xception(seg) (Ours) | 95.15 | **87.61** | 92.08 | **72.82** | 57.58 | 81.04 |
| Xception(fusion) (Ours) | **95.28** | 86.96 | **93.24** | 71.66 | 63.80 | **82.19** |

## D. Manipulation Localization

In this subsection, we evaluate the manipulation localization performance of the proposed model both quantitatively and qualitatively.

*1) Quantitative evaluation:* As we apply the multi-scale feature learning strategy in the proposed framework, three predicted manipulation maps can be obtained for each given face. Herein, we apply a weighted summation strategy over the three predicted maps to determine the final manipulation localization map:

$$\hat{M}_f = \gamma_1 \hat{M}_1^{\uparrow} + \gamma_2 \hat{M}_2^{\uparrow} + \gamma_3 \hat{M}_3^{\uparrow}, \tag{9}$$

where $\hat{M}_f$ represents the final determined manipulation localization map. The parameters $\gamma_1$, $\gamma_2$, and $\gamma_3$ are used to weigh the predicted maps. We resize the predicted masks $\hat{M}_1$, $\hat{M}_2$, and $\hat{M}_3$ ($\hat{M}_1$, $\hat{M}_2$, and $\hat{M}_3$ are specified in Eq. (3)) to the original input face size $299 \times 299$, obtaining the aligned manipulation localization maps $\hat{M}_1^{\uparrow}$, $\hat{M}_2^{\uparrow}$, and $\hat{M}_3^{\uparrow}$.

Intuitively, the features extracted from the deep layer tend to carry rich semantic information. Thus we allocate relatively larger weight to $\hat{M}_3^{\uparrow}$. To study the best weighting strategy, we report the Intersection over Union (IoU) score for both
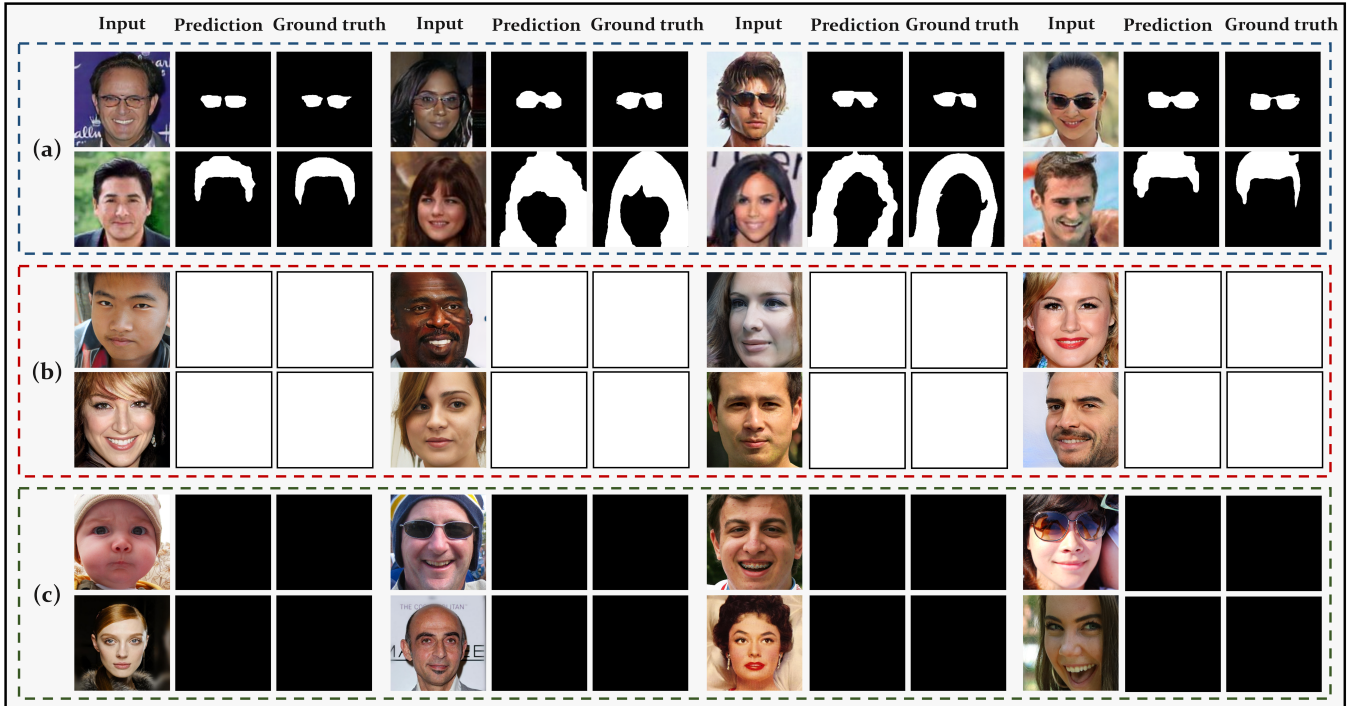
Fig. 10. Face manipulation localization results on our collected dataset. (a). attribute manipulated faces; (b). entirely synthetic faces; (c). real faces.

TABLE VII
FACE MANIPULATION LOCALIZATION ACCURACY (IoU) FOR
LOW-QUALITY (LQ) AND HIGH-QUALITY (HQ) FACES WITH DIFFERENT
WEIGHTING STRATEGIES.

| $\gamma_1$ | $\gamma_2$ | $\gamma_3$ | IoU (LQ) | IoU (HQ) |
|------|------|------|----------|----------|
| 0.0 | 0.0 | 1.0 | 0.8399 | 0.9277 |
| 0.0 | 0.1 | 0.9 | 0.8408 | 0.9291 |
| 0.1 | 0.1 | 0.8 | 0.8406 | 0.9304 |
| 0.1 | 0.2 | 0.7 | **0.8413** | **0.9305** |
| 0.1 | 0.3 | 0.6 | 0.8371 | 0.9240 |
| 0.1 | 0.4 | 0.5 | 0.8270 | 0.9110 |
| 0.2 | 0.2 | 0.6 | 0.8373 | 0.9247 |
| 0.2 | 0.3 | 0.5 | 0.8313 | 0.9177 |
| 0.2 | 0.4 | 0.4 | 0.8205 | 0.9015 |

high-quality and low-quality faces in Table VII. Comparing the IoUs in the first row ($\gamma_1$=0.0, $\gamma_2$=0.0, $\gamma_3$=1.0) and the IoUs in the fourth row ($\gamma_1$=0.1, $\gamma_2$=0.2, $\gamma_3$=0.7), we can conclude that the multi-scale feature learning strategy benefits the manipulation localization performance because leveraging the information of all three features leads to a better localization accuracy. As such, the values of $\gamma_1$, $\gamma_2$, and $\gamma_3$ are set as 0.1, 0.2, and 0.7. Our model can achieve remarkable 0.8413 and 0.9305 IoU localization performance for low-quality and high-quality faces, respectively, demonstrating the effectiveness of the designed semantic map prediction module.

*2) Qualitative evaluation on FF++ dataset:* To further study the manipulation localization performance of the proposed model, we qualitatively present the results of face manipulation localization on high-quality faces and low-quality

faces in Fig. 9 (a) and Fig. 9 (b). The faces in different rows represent the fake faces created by corresponding manipulation methods and pristine real faces. The input face, predicted localization map and the corresponding ground truth mask are shown in the left, middle, and right columns. Comparing with the ground-truth masks, it can be seen that the manipulated regions can be well captured for both low-quality and high-quality faces by the proposed model. Furthermore, our model can accurately determine the forgery locations for the faces with various head-poses and very poor lighting conditions (e.g., NT 1st example), which further validates the robustness of the proposed model.

*3) Qualitative evaluation on our collected dataset:* In addition to the FF++ dataset, which manipulates faces in the whole face region, we further consider two pervasive face manipulation types in this subsection: entirely synthetic faces and attribute manipulated faces. Herein, we employ the popular StyleGAN [57] and PGGAN [58] to produce 40000 entirely synthetic faces and use StarGAN [59] and AttGAN [60] to generate 40000 attribute-manipulated faces (two attribute manipulations are considered in this paper: glasses and hairs). We further collect 40000 real faces from the FFHQ [57] and CelebA-HQ [58] datasets. The training, validation, and test sets are split as 8:1:1. We denote the collected Face Manipulation Localization Dataset as FMLD. The manipulation localization performance of the proposed model on FMLD is shown in Fig. 10. Again, the proposed model achieves outstanding forgery localization performance.

*4) Comparison with SOTA forgery localization methods:* In this part, we compare the proposed model with other methods using masks in the training stage. We respectively report the manipulation detection and localization results in Table VIII

TABLE VIII
DETECTION PERFORMANCE ON FF++ LQ (QP=40) DATASET COMPARED
WITH THE METHODS USING MASKS.

| Methods | AUC(%) ↑ | ACC(%) ↑ | AP(%) ↑ | FPR(%) ↓ | FNR(%) ↓ | EER(%) ↓ |
|---|---|---|---|---|---|---|
| Multitask | 75.73 | 75.10 | 91.76 | **43.77** | 20.18 | 30.79 |
| DFFD Reg. | <u>84.95</u> | <u>83.70</u> | <u>95.24</u> | 48.78 | **8.18** | <u>23.03</u> |
| DFFD Mam. | 83.87 | 83.63 | 94.73 | 46.06 | 8.95 | 23.32 |
| Ours | **86.87** | **84.56** | **95.98** | <u>44.18</u> | <u>8.26</u> | **21.03** |

TABLE IX
MANIPULATION LOCALIZATION PERFORMANCE COMPARED WITH
STATE-OF-THE-ARTS LOCALIZATION METHODS.

| Dataset | FF++ HQ | | FMLD | |
|---|---|---|---|---|
| Methods | PBCA(%) ↑ | IINC(%) ↓ | PBCA(%) ↑ | IINC(%) ↓ |
| Multitask [55] | 94.88 | 4.46 | 98.59 | 3.52 |
| DFFD Reg. [12] | 94.85 | 4.57 | 98.72 | 3.31 |
| DFFD Mam. [12] | 91.45 | 13.09 | 96.86 | 23.93 |
| Ours | **95.77** | **3.62** | **99.06** | **2.53** |

and Table IX. It is worth mentioning that noise information is not used in our method in these two evaluations for fair comparisons. We present the detection performance on FF++ LQ (QP=40) dataset in Table VIII. Multitask [55] designs a U-shape CNN that simultaneously detects manipulated faces and locates the manipulated regions for each query. DFFD [12] uses mask attention mechanism to perform face manipulation detection and localization, and DFFD Reg. and DFFD Mam. are two different attention mechanisms of DFFD. In this table, we use the bold font to highlight the best results while underlining the second-best results among all listed methods. It can be observed that our method achieves the best AUC, ACC, AP, and EER scores and the second-best FPR and FNR performance. These performance boosts benefit from the proposed multi-scale mask supervision. Moreover, we further report the manipulation localization results on the FF++ HQ (QP=23) and FMLD datasets of each method in Table IX. Following DFFD, we use Pixel-wise Binary Classification Accuracy (PBCA) and Inverse Intersection Non-Containment (IINC) as the evaluation metrics. In Table IX, the proposed method achieves the best localization performance on both FF++ HQ and FMLD datasets. We further present localization examples of these methods in Fig. 11. Compared with Multitask, our method has a superior pixel-wise localization accuracy. Compared with DFFD, the proposed method has a more precise localization performance on the boundary regions, which mainly benefit from our multi-scale learning strategy.

*5) Discussion:* In this subsection, we quantitatively and qualitatively evaluate the manipulation localization performance. The proposed model captures the high-level semantic information from the multi-scale features, with state-of-the-art localization accuracy and segmentation performance for FF++ faces, attribute-manipulated faces, and entirely synthetic faces.

*E. Ablation Study*

To validate the effectiveness of the semantic map prediction module, the noise map prediction module, and the multi-scale feature extraction strategy in the proposed framework, we conduct extensive ablation experiments in this subsection.

TABLE X
AUC SCORES OF THE XCEPTION BACKBONE EQUIPPED WITH SEMANTIC
MAP PREDICTION MODULE AND NOISE MAP PREDICTION MODULE ON
LOW-QUALITY FACES.

| Xception | Noise | Mask | AUC(%) ↑ | AP(%) ↑ | EER(%) ↓ |
|---|---|---|---|---|---|
| √ | - | - | 84.38 | 94.63 | 22.97 |
| √ | √ | - | 82.15 | 94.40 | 24.84 |
| √ | - | √ | 86.87 | 95.98 | 21.03 |
| √ | √ | √ | **87.10** | **96.18** | **20.94** |

*1) Effectiveness of sub-modules:* To study the effectiveness of the semantic map prediction module and the noise map prediction module, we report the manipulation detection results of the Xception backbone equipped with different modules on low-quality faces in Table X. By comparing with the AUC, AP and EER scores in the third row and the detection result of the Xception baseline, we can conclude that the semantic map prediction module leads the model to mine more artifact clues and improve the detection performance. On the other hand, the detection performance drops slightly when equipping the noise prediction module on the Xception backbone. The reason is that forcing the model to only focus on the noise-level clues may cause severe image information loss, resulting in detection performance drops. In this work, the noise clue is regarded as the auxiliary information that plays a complementary role to the semantic map prediction module.

Comparing with the detection results in the third row, the detection performance in the last row gains specific improvement, proving the effectiveness of the noise map prediction module. To better validate the effectiveness of each module, we also show the ROC curves in Fig. 12. We can observe that the proposed model achieves the best AUC performance and the best TPR at a lower FPR value, while low FPR is one of the most challenging face manipulation detection scenarios [9].

We can be concluded that the semantic map prediction module and the noise map prediction module indeed help improve the detection capability of the proposed model.

*2) Effectiveness of multi-scale feature learning:* We further conduct an ablation study to demonstrate the effectiveness of the multi-scale feature learning strategy. The manipulation detection performance of different feature extraction strategies is listed in Table XI. As illustrated in Fig. 2, $f_1$, $f_2$, and $f_3$ represent the features extracted from the shallow, middle, and deep layers of the Xception backbone, and each feature layer will be processed by the semantic map prediction module and the noise map prediction module.

To better illustrate the effectiveness of each feature layer, we also show the ROC curves of different feature extraction strategies in Fig. 13. From the results in Table XI and Fig. 13, we can conclude that 1) the usage of deep feature layer $f_3$ leads to the most AUC score gain as it contains much semantic information; and 2) each extracted feature layer contributes the final detection performance improvements, verifying the effectiveness of the proposed multi-scale feature learning strategy.
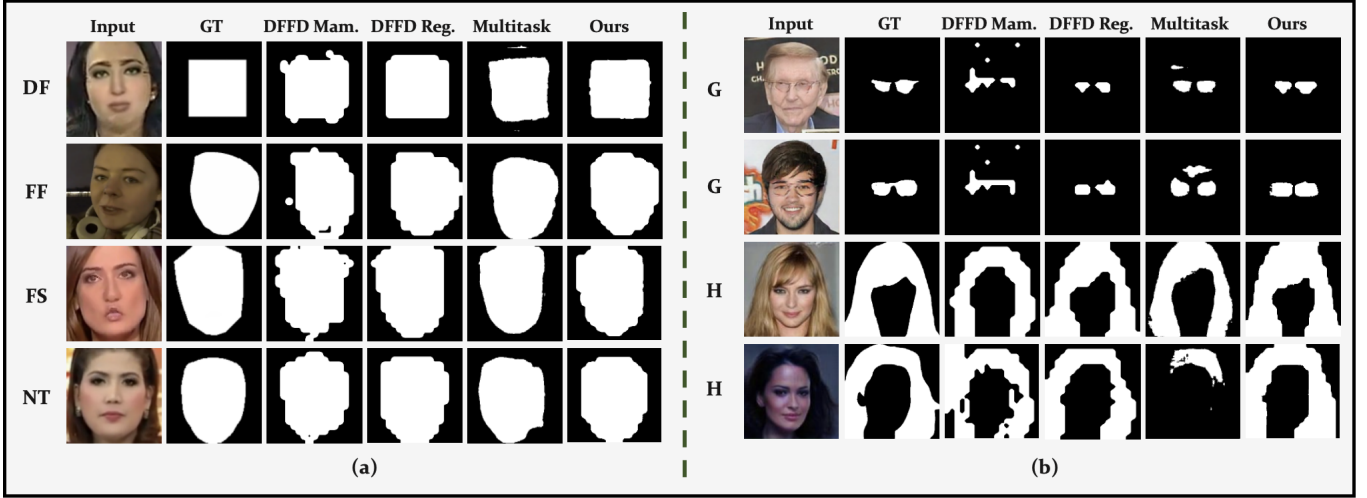
Fig. 11. Face manipulation localization results of different methods on (a) FF++ HQ dataset and (b) FMLD. DF, FF, FS, and NT represent Deepfakes, Face2Face, FaceSwap, and NeuralTextures. G and H indicate two attribute manipulations: glasses and hair.

TABLE XI
ABLATION STUDY OF THE PROPOSED MODEL WITH DIFFERENT FEATURE LAYER EXTRACTION STRATEGIES ON LOW-QUALITY FACES.

| Xception | $f_3$ | $f_2$ | $f_1$ | AUC(%) ↑ | AP(%) ↑ | EER(%) ↓ |
|---|---|---|---|---|---|---|
| √ | - | - | - | 84.38 | 94.63 | 22.97 |
| √ | √ | - | - | 86.39 | 95.87 | 22.02 |
| √ | √ | √ | - | 86.76 | **96.20** | 21.27 |
| √ | √ | √ | √ | **87.10** | 96.18 | **20.94** |
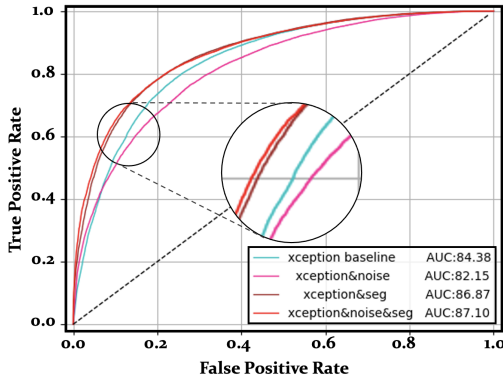


Fig. 12. Results for the Xception baseline and the three proposed models (ROC).



Fig. 13. Results for the Xception baseline and the proposed models with three different feature extraction strategies (ROC).

*3) Study of feature aggregation strategy:* Thanks to the proposed multi-scale learning strategy, the three features extracted from the backbone should contain abundant information. In this part, we further study the impact of the feature aggregation on the forgery classification performance. The illustration of the feature aggregation strategy is shown in Fig. 14. The $f_1$, $f_2$, and $f_3$ are synchronously fed to semantic map prediction module and noise map prediction module. Sequentially, we feed forward $f_1$ and $f_2$ into Size Align Blocks (SAB) that consist of several convolution and pooling layers for feature size alignment. The aligned noise and mask features are concatenated along the channel dimension. Then, we conduct the f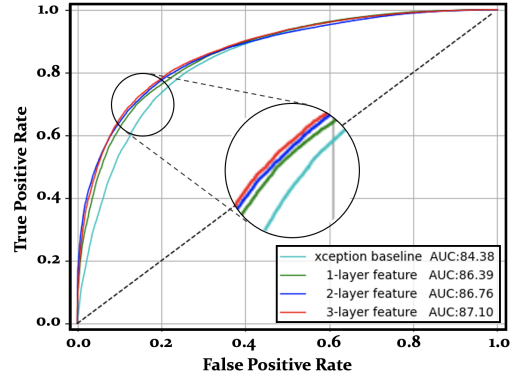eature spatial attention by multiplying the mask map predicted by the last $f_3$-based semantic map prediction module with the fused feature. The attended feature is finally sent to the classification module to identify the authenticity of the input query face. We summarize the feature aggregation process as follows,

$$F_{att} = \{\underbrace{[SAB_1^m(f_1^m) \oplus SAB_2^m(f_2^m) \oplus f_3^m]}_{F_{cat}^m}$$
$$\oplus \underbrace{[SAB_1^n(f_1^n) \oplus SAB_2^n(f_2^n) \oplus f_3^n]}_{F_{cat}^n}\} \otimes M_3 \quad (10)$$

where the superscripts $m$ and $n$ indicate the mask and noise prediction modules. $\oplus$ indicates the concatenation operation. $M_3$ is the predicted mask from the semantic map prediction module, and $\otimes$ denotes the spatial attention process.

We report the detection results of such feature aggregation strategy in Table XII (denoted as Aggregation), and the results of using single feature $f_3$ (denoted as Single $f_3$) are also listed for reference. We can observe that the feature aggregation indeed improves the detection performance. It is worth

TABLE XII
MANIPULATION DETECTION RESULTS COMPARED WITH FEATURE
AGGREGATION STRATEGY.

| Method | AUC(%) ↑ | ACC(%) ↑ | AP(%) ↑ | EER(%) ↓ |
|---|---|---|---|---|
| Aggregation | **87.36** | **85.02** | **96.24** | **20.75** |
| Single $f_3$ | 87.10 | 84.84 | 96.18 | 20.94 |

TABLE XIII
DETECTION PERFORMANCE OF THE PROPOSED TWO-STREAM
MULTI-SCALE FRAMEWORK ON VGG16 AND EFFICIENTNET BACKBONE.

| Methods | AUC(%) ↑ | ACC(%) ↑ | AP(%) ↑ | EER(%) ↓ |
|---|---|---|---|---|
| VGG16 [61] | 81.45 | 82.46 | 93.41 | 25.49 |
| VGG16 & Mask. | 82.70 | 82.53 | 94.61 | 24.99 |
| VGG16 & Noise | 79.14 | 79.78 | 92.84 | 27.78 |
| VGG16 & Fusion | **83.45** | **83.27** | **94.62** | **24.25** |
| EfficientNet | 84.19 | 84.05 | 94.30 | 22.88 |
| EfficientNet & Mask. | 87.25 | 84.70 | 96.37 | 20.64 |
| EfficientNet & Noise | 84.08 | 83.97 | 94.01 | 23.02 |
| EfficientNet & Fusion | **87.96** | **85.28** | **96.39** | **19.63** |

TABLE XIV
ABLATION STUDY FOR THE EFFECTS OF LOSS WEIGHTS $\lambda_1$ AND $\lambda_2$.

| $\lambda_2$ | $\lambda_1$ | AUC(%) | AP(%) | EER(%) |
|---|---|---|---|---|
| 0.0 | - | 84.38 | 94.63 | 22.97 |
| 0.1 | - | 84.91 | 95.33 | 23.34 |
| 0.5 | - | 86.35 | 95.63 | 21.72 |
| 1.0 | - | 86.87 | 95.98 | 21.03 |
| 5.0 | - | 86.76 | 96.09 | 21.32 |
| 10.0 | - | 86.72 | 95.90 | 21.24 |
| 1.0 | √ | **87.10** | **96.18** | **20.94** |

TABLE XV
DETECTION PERFORMANCE USING DIFFERENT NOISE FILTERS.

| Method | AUC(%) ↑ | ACC(%) ↑ | AP(%) ↑ | FPR(%) ↓ | FNR(%) ↓ | EER(%) ↓ |
|---|---|---|---|---|---|---|
| Noise input | 82.07 | 81.14 | 94.42 | 47.79 | 11.63 | 25.71 |
| SRM HP [47] | 81.38 | 81.95 | 93.67 | 51.61 | 9.66 | 26.19 |
| CNN-based [64] | 82.66 | 82.63 | 94.40 | 49.74 | 9.27 | 24.84 |
| Ours | 82.15 | 82.75 | 94.55 | 53.63 | 8.15 | 25.66 |

mentioning that this also makes the model bigger and requires more computational resources.

*4) Evaluation of the proposed learning scheme on other backbones:* Extensive experiments have demonstrated the effectiveness of the proposed framework on the Xception backbone. In this part, we further apply the proposed two-stream multi-scale learning architecture on the VGG16 [61] and EfficientNet [56] backbones to verify that our framework can be conveniently transferred to other detectors in a plug-and-play manner. The detection results are reported in Table XIII. Unsurprisingly, the semantic-level supervision boosts the detection accuracy. And the noise fingerprint playing a complementary role further empowers the final classification performance.

*5) Effectiveness of noise features:* In the training and test stages, we exploit the noise feature rather than directly use the noise map to conduct the binary classification. Directly feeding the noise maps into the model will primarily suppress the image content and may lead the model to make wrong decisions, as noise is a weak signal and can be easily overwhelmed by imperfectly removed image content [62]. This work exploits noise fingerprints more flexibly, extracting the informative noise features from input faces and fusing them with high-level semantic features. Besides, noise-based fingerprints can emphasize high-frequency clues and expose more spatial local artifacts, which are largely ignored in the high-level semantic feature. Therefore, the noise feature plays a complementary role to the high-level semantic feature, and the two-stream learning strategy can be conducted by mining high-level and low-level artifacts in the feature space.

As defined in Eq. (1), the objective function of overall framework consists of the following three components: the classification loss $L_c$, noise prediction loss $L_n$, and binary mask prediction loss $L_b$, and $\lambda_1$ and $\lambda_2$ are hyper-parameters to weigh the loss components. To better demonstrate the effectiveness of the noise prediction module, we conduct an ablation study for the effects of loss weights $\lambda_1$ and $\lambda_2$. We

report the detection results in Table. XIV. We can see that when the overall framework is trained without noise loss component (*i.e.*, $\lambda_1 = 0$), no matter how we fine-tune the value of $\lambda_2$ from 0.1 to 10.0, the manipulation detection performances are consistently lower than the model trained with the noise loss component $L_n$. This experiment indicates that the adoption of noise information does improve the framework's face manipulation detection performance.

To further demonstrate the effectiveness of devised noise map prediction module, we visualize the predicted noise maps, input faces, and ground-truth noise maps in Fig. 15. The top row shows the input faces. The middle and bottom rows present the corresponding ground-truth and predicted noise maps. It can be observed that the noise predictor can efficiently predict the noise maps, demonstrating that our model has learned the low-level noise information. We further use red boxes to mark the local artifacts in the predicted noise maps, and these clues are hard to be recognized in the RGB space. We can conclude that, in this model, the noise features carrying informative low-level local clues have been successfully extracted, thus leading to a certain face manipulation detection performance gain.

*6) Effects of different noise filters:* To further explore the impacts of other noise filters, we conduct more experiments using the SRM high-pass filter and the CNN-based filter. Table XV also reports the detection results of directly taking the noise map as input for reference. SRM has been demonstrated its effectiveness in image forensics tasks [18], [63] as it can expose the high-frequency artifacts in the manipulated images. Moreover, the recently proposed CNN-based filter [64] has been widely used in various image denoising tasks. Compared with results of directing use noise as input, our framework exploiting the noise fingerprints in the feature space has a better detection performance, as shown in the last row in Table XV. Furthermore, the results of using the three different noise filters achieve competitive forgery detection performances.

*7) Discussion:* Herein, we conduct a series of ablation studies to evaluate the effectiveness of the dedicated learn-
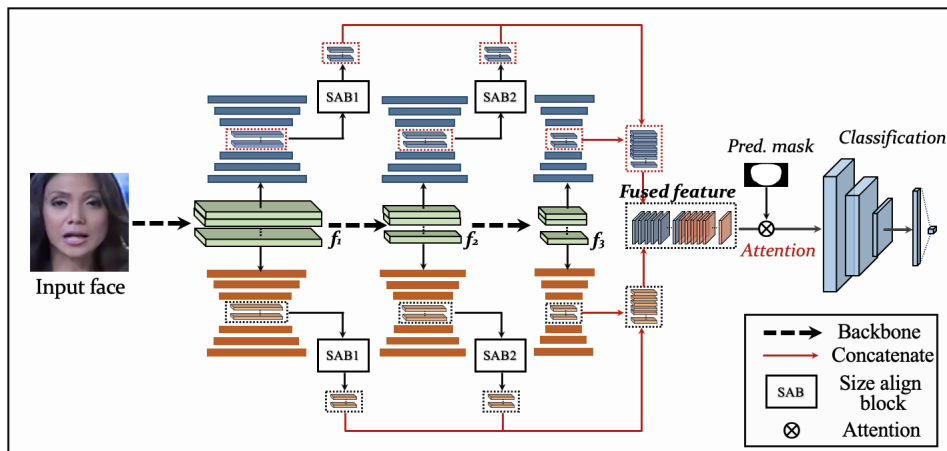
Fig. 14. Illustration of feature aggregation strategy. The $f_1$, $f_2$, and $f_3$ are synchronously fed to semantic map prediction module and noise map prediction module. Sequentially, we feed forward $f_1$ and $f_2$ into Size Align Blocks (SAB) that consist of several convolution and pooling layers for feature size alignment. The aligned noise and mask features are concatenated along the channel dimension. Then, we conduct the feature spatial attention by multiplying the mask map predicted by the last $f_3$-based semantic map prediction module with the fused feature. And the attended feature is finally sent to the classification module to identify the authenticity of input query face.
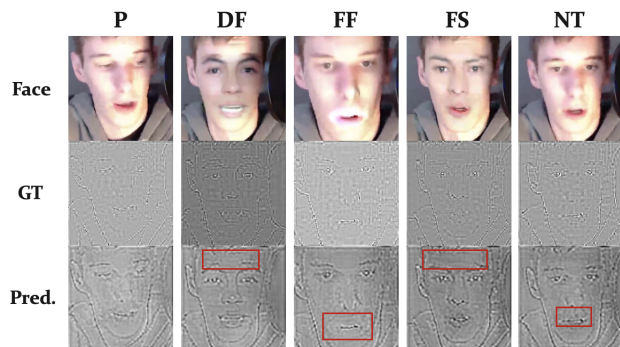


Fig. 15. Visualization of the noise map predictions. The top row shows the input faces. The middle and bottom rows respectively present the ground-truth and predicted noise maps (The noise map intensity has been enlarged for better visualization). P: Pristine, DF: Deepfakes, FF: Face2Face, FS: FaceSwap, NT: NeuralTextures.

ing strategy and framework architecture. We also study the impacts of feature aggregation strategy, different backbones, noise features, and different noise filters to comprehensively interpret the proposed framework. The qualitative and quantitative experimental results demonstrate that the designed submodules and the multi-scale learning strategy lead the model to mine more artifact clues from the input faces, which further improves the final forgery detection performance.

## V. CONCLUSIONS AND FUTURE WORK

This paper presented a novel framework to tackle the face manipulation problem. In particular, we introduced two complementary tasks, including semantic map prediction and noise map prediction, to capture the semantic-level and noise-level information to perform both face manipulation detection and forgery localization. Extensive experimental results show that the proposed two-stream multi-scale framework outperforms the state-of-the-art detection methods and state-of-the-art cross-dataset detection methods for both high-quality and low-quality faces.

The proposed semantic map prediction module enables the model to perform face manipulation localization, and it also constrains the model to focus on manipulated regions, thus leading to a better binary and multi-class classification performance. Furthermore, the noise map prediction module serves as a complementary module, and it provides significative noise-level clues and subsequently empowers the final decision-making.

Last but not least, an ablation study demonstrated the effectiveness of the semantic map prediction module and the noise map prediction module, and the multi-scale feature learning strategy indeed helps the model improve its manipulation detection performance.

While our proposed method has shown to be effective for face manipulation detection and localization tasks, we limit our scope to only four FF++ attacks (e.g., DeepFakes, Face2Face, FaceSwap, NeuralTextures) and two GAN attacks (attribute-manipulated faces and entirely synthetic faces). Therefore, adapting our method to unseen face manipulation is worth investigating in the future, although we already provided a good indication of the method's performance considering the complex cross-dataset validation. On the other hand, since natural videos also face a similar threat of malicious manipulation, it is also worth applying our method to the manipulation detection based on other types of multimedia content in the future.

## VI. ACKNOWLEDGEMENT

## REFERENCES

[1] "Github deepfake faceswap," https://github.com/deepfakes/faceswap, 2018.

[2] "Deepfakes faceswap," [EB/OL], 2019, https://github.com/deepfakes/faceswap.

[3] J. Thies, M. Zollhofer, M. Stamminger, C. Theobalt, and M. Nießner, "Face2face: Real-time face capture and reenactment of rgb videos," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 2387–2395.

[4] J. Thies, M. Zollhöfer, and M. Nießner, "Deferred neural rendering: Image synthesis using neural textures," *ACM Transactions on Graphics (TOG)*, vol. 38, no. 4, pp. 1–12, 2019.

[5] "Fakeapp," [EB/OL], 2018, https://www.malavida.com/en/soft.

[6] Y. Li, M.-C. Chang, and S. Lyu, "In ictu oculi: Exposing ai created fake videos by detecting eye blinking," in *2018 IEEE International Workshop on Information Forensics and Security (WIFS)*. IEEE, 2018, pp. 1–7.

[7] X. Yang, Y. Li, and S. Lyu, "Exposing deep fakes using inconsistent head poses," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 8261–8265.

[8] L. Li, J. Bao, T. Zhang, H. Yang, D. Chen, F. Wen, and B. Guo, "Face x-ray for more general face forgery detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 5001–5010.

[9] Y. Qian, G. Yin, L. Sheng, Z. Chen, and J. Shao, "Thinking in frequency: Face forgery detection by mining frequency-aware clues," in *European Conference on Computer Vision*. Springer, 2020, pp. 86–103.

[10] H. Liu, X. Li, W. Zhou, Y. He, H. Xue, W. Zhang, and N. Yu, "Spatial-phase shallow learning: rethinking face forgery detection in frequency domain," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 772–781.

[11] I. Masi, A. Killekar, R. M. Mascarenhas, S. P. Gurudatt, and W. AbdAlmageed, "Two-branch recurrent network for isolating deepfakes in videos," in *European Conference on Computer Vision*. Springer, 2020, pp. 667–684.

[12] H. Dang, F. Liu, J. Stehouwer, X. Liu, and A. K. Jain, "On the detection of digital face manipulation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 5781–5790.

[13] M. Du, S. Pentyala, Y. Li, and X. Hu, "Towards generalizable deepfake detection with locality-aware autoencoder," in *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, 2020, pp. 325–334.

[14] A. Chintha, B. Thai, S. J. Sohrawardi, A. Bhatt, A. Hickerson, M. Wright, and R. Ptucha, "Recurrent convolutional structures for audio spoof and video deepfake detection," *IEEE Journal of Selected Topics in Signal Processing*, vol. 14, no. 5, pp. 1024–1037, 2020.

[15] K. Songsri-in and S. Zafeiriou, "Complement face forensic detection and localization with faciallandmarks," *arXiv preprint arXiv:1910.05455*, 2019.

[16] X. Li, Y. Lang, Y. Chen, X. Mao, Y. He, S. Wang, H. Xue, and Q. Lu, "Sharp multiple instance learning for deepfake video detection," in *Proceedings of the 28th ACM International Conference on Multimedia*, 2020, pp. 1864–1872.

[17] J. H. Bappy, A. K. Roy-Chowdhury, J. Bunk, L. Nataraj, and B. Manjunath, "Exploiting spatial structure for localizing manipulated image regions," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 4970–4979.

[18] P. Zhou, X. Han, V. I. Morariu, and L. S. Davis, "Learning rich features for image manipulation detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 1053–1061.

[19] P. Zhou, B.-C. Chen, X. Han, M. Najibi, A. Shrivastava, S.-N. Lim, and L. Davis, "Generate, segment, and refine: Towards generic manipulation segmentation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 07, 2020, pp. 13 058–13 065.

[20] A. George and S. Marcel, "Deep pixel-wise binary supervision for face presentation attack detection," in *2019 International Conference on Biometrics (ICB)*. IEEE, 2019, pp. 1–8.

[21] Y. Liu, J. Stehouwer, A. Jourabloo, and X. Liu, "Deep tree learning for zero-shot face anti-spoofing," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4680–4689.

[22] W. Sun, Y. Song, C. Chen, J. Huang, and A. C. Kot, "Face spoofing detection based on local ternary label supervision in fully convolutional networks," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 3181–3196, 2020.

[23] A. C. Popescu and H. Farid, "Statistical tools for digital forensics," in *international workshop on information hiding*. Springer, 2004, pp. 128–147.

[24] B. Mahdian and S. Saic, "Using noise inconsistencies for blind image forensics," *Image and Vision Computing*, vol. 27, no. 10, pp. 1497–1503, 2009.

[25] S. Lyu, X. Pan, and X. Zhang, "Exposing region splicing forgeries with blind local noise estimation," *International journal of computer vision*, vol. 110, no. 2, pp. 202–221, 2014.

[26] J. Fridrich, J. Lukas, and M. Goljan, "Digital camera identification from sensor noise," *IEEE Transactions on Information Security and Forensics*, vol. 1, no. 2, pp. 205–214, 2006.

[27] P. Korus and J. Huang, "Multi-scale analysis strategies in prnu-based tampering localization," *IEEE Transactions on Information Forensics and Security*, vol. 12, no. 4, pp. 809–824, 2016.

[28] Y. Quan and C.-T. Li, "On addressing the impact of iso speed upon prnu and forgery detection," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 190–202, 2020.

[29] D. Cozzolino and L. Verdoliva, "Noiseprint: A cnn-based camera model fingerprint," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 144–159, 2019.

[30] S. Zagoruyko and N. Komodakis, "Learning to compare image patches via convolutional neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 4353–4361.

[31] A. Jourabloo, Y. Liu, and X. Liu, "Face de-spoofing: Anti-spoofing via noise modeling," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 290–306.

[32] Y. Ren, Y. Hu, B. Liu, Y. Xie, and Y. Wang, "Face anti-spoofing with a noise-attention network using color-channel difference images," in *International Conference on Artificial Neural Networks*. Springer, 2020, pp. 517–526.

[33] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1251–1258.

[34] A. Rossler, D. Cozzolino, L. Verdoliva, C. Riess, J. Thies, and M. Nießner, "Faceforensics++: Learning to detect manipulated facial images," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 1–11.

[35] J. Lukas, J. Fridrich, and M. Goljan, "Digital camera identification from sensor pattern noise," *IEEE Transactions on Information Forensics and Security*, vol. 1, no. 2, pp. 205–214, 2006.

[36] M. K. Mihcak, I. Kozintsev, and K. Ramchandran, "Spatially adaptive statistical modeling of wavelet image coefficients and its application to denoising," in *1999 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings. ICASSP99 (Cat. No. 99CH36258)*, vol. 6. IEEE, 1999, pp. 3253–3256.

[37] Y. Li, X. Yang, P. Sun, H. Qi, and S. Lyu, "Celeb-df: A large-scale challenging dataset for deepfake forensics," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 3207–3216.

[38] "Deepfakedetection," https://ai.googleblog.com/2019/09/contributing-data-to-deepfake-detection.html, 2020.

[39] B. Dolhansky, J. Bitton, B. Pflaum, J. Lu, R. Howes, M. Wang, and C. Canton Ferrer, "The deepfake detection challenge dataset," *arXiv e-prints*, pp. arXiv–2006, 2020.

[40] D. E. King, "Dlib-ml: A machine learning toolkit," *The Journal of Machine Learning Research*, vol. 10, pp. 1755–1758, 2009.

[41] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "Pytorch: An imperative style, high-performance deep learning library." in *NeurIPS*, 2019.

[42] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[43] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.

[44] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3431–3440.

[45] H. Noh, S. Hong, and B. Han, "Learning deconvolution network for semantic segmentation," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1520–1528.

[46] Y. Nirkin, L. Wolf, Y. Keller, and T. Hassner, "Deepfake detection based on discrepancies between faces and their context," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.

[47] J. Fridrich and J. Kodovsky, "Rich models for steganalysis of digital images," *IEEE Transactions on Information Forensics and Security*, vol. 7, no. 3, pp. 868–882, 2012.

[48] D. Cozzolino, G. Poggi, and L. Verdoliva, "Recasting residual-based local descriptors as convolutional neural networks: an application to image forgery detection," in *Proceedings of the 5th ACM Workshop on Information Hiding and Multimedia Security*, 2017, pp. 159–164.

[49] B. Bayar and M. C. Stamm, "A deep learning approach to universal image manipulation detection using a new convolutional layer," in *Proceedings of the 4th ACM Workshop on Information Hiding and Multimedia Security*, 2016, pp. 5–10.

[50] N. Rahmouni, V. Nozick, J. Yamagishi, and I. Echizen, "Distinguishing computer graphics from natural images using convolution neural networks," in *2017 IEEE Workshop on Information Forensics and Security (WIFS)*. IEEE, 2017, pp. 1–6.

[51] D. Afchar, V. Nozick, J. Yamagishi, and I. Echizen, "Mesonet: a compact facial video forgery detection network," in *2018 IEEE International Workshop on Information Forensics and Security (WIFS)*. IEEE, 2018, pp. 1–7.

[52] H. Zhao, W. Zhou, D. Chen, T. Wei, W. Zhang, and N. Yu, "Multi-attentional deepfake detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 2185–2194.

[53] P. Korshunov and S. Marcel, "Deepfakes: a new threat to face recognition? assessment and detection," *arXiv preprint arXiv:1812.08685*, 2018.

[54] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[55] H. H. Nguyen, F. Fang, J. Yamagishi, and I. Echizen, "Multi-task learning for detecting and segmenting manipulated facial images and videos," in *2019 IEEE 10th International Conference on Biometrics Theory, Applications and Systems (BTAS)*. IEEE, 2019, pp. 1–8.

[56] M. Tan and Q. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *International Conference on Machine Learning*. PMLR, 2019, pp. 6105–6114.

[57] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4401–4410.

[58] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of gans for improved quality, stability, and variation," in *International Conference on Learning Representations*, 2018.

[59] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo, "Stargan: Unified generative adversarial networks for multi-domain image-to-image translation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 8789–8797.

[60] Z. He, W. Zuo, M. Kan, S. Shan, and X. Chen, "Attgan: Facial attribute editing by only changing what you want," *IEEE transactions on image processing*, vol. 28, no. 11, pp. 5464–5478, 2019.

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

[62] L. Verdoliva, "Media forensics and deepfakes: an overview," *IEEE Journal of Selected Topics in Signal Processing*, vol. 14, no. 5, pp. 910–932, 2020.

[63] Y. Wu, W. AbdAlmageed, and P. Natarajan, "Mantra-net: Manipulation tracing network for detection and localization of image forgeries with anomalous features," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 9543–9552.

[64] C. Tian, Y. Xu, Z. Li, W. Zuo, L. Fei, and H. Liu, "Attention-guided cnn for image denoising," *Neural Networks*, vol. 124, pp. 117–129, 2020.