

FUSING GLOBAL AND LOCAL FEATURES FOR GENERALIZED AI-SYNTHESIZED IMAGE DETECTION

Yan Ju[†], Shan Jia[†], Lipeng Ke[†], Hongfei Xue[†], Koki Nagano[‡], and Siwei Lyu[†]

[†]University at Buffalo, State University of New York
[‡]NVIDIA

ABSTRACT

With the development of the Generative Adversarial Networks (GANs) and DeepFakes, AI-synthesized images are now of such high quality that humans can hardly distinguish them from real images. It is imperative for media forensics to develop detectors to expose them accurately. Existing detection methods have shown high performance in generated images detection, but they tend to generalize poorly in the real-world scenarios, where the synthetic images are usually generated with unseen models using unknown source data. In this work, we emphasize the importance of combining information from the whole image and informative patches in improving the generalization ability of AI-synthesized image detection. Specifically, we design a two-branch model to combine global spatial information from the whole image and local informative features from multiple patches selected by a novel patch selection module. Multi-head attention mechanism is further utilized to fuse the global and local features. We collect a highly diverse dataset synthesized by 19 models with various objects and resolutions to evaluate our model. Experimental results demonstrate the high accuracy and good generalization ability of our method in detecting generated images. Our code is available at <https://github.com/littlejuyan/FusingGlobalandLocal>.

Index Terms— AI-synthesized Image Detection, Image Forensics, Feature Fusion, Attention Mechanism.

1. INTRODUCTION

Recently, the creation of realistic synthetic media has seen rapid developments with ever improving visual qualities and run time efficiencies. Notable examples of AI-synthesized media include the high-quality images synthesized by the Generative Adversarial Networks (GANs) [1], and videos with face swaps or puppetry created by the auto-encoder models [2]. The synthetic images/videos have become challenging for humans to distinguish [3], and correspondingly, a slew of detection methods [4, 5] have been developed to mitigate the potential risks posed by such AI-synthesized images.

Existing detection methods generally employ deep neural networks to automatically extract discriminative features, and then feed them into a binary classifier to tell real or fake. Most




Fig. 1. Detectors trained with high resolution frontal face images generated by the seen model have high accuracies. What about the images with various resolutions, objects and manipulation types from unseen models?

of the methods focus on global features extracted from the entire images [6, 7, 8, 9]. Inspired by the success in fine-grained classification field, a recent work [7] brings a novel perspective by taking the binary image detection as a fine-grained classification task, and proposes to capture local features from multiple face attentive regions.

Current detectors have provided an excellent performance in a closed (*i.e.*, intra-dataset) scenario, where the testing and training images are generated by the same models trained on similar source data. However, their performance tends to decrease considerably in the open-world applications, where the testing images come from a different domain, and generated by unknown models or from unseen source data. For example, [10] evaluates several SOTA detectors under the real-world conditions, including cross-model, cross-data, and post-processing, and concludes that current detection models are not yet robust enough for real applications.

To improve the detection generalization, several recent works propose to utilize new deep learning techniques, such as transfer learning [11], incremental learning [12], contrastive learning [13], representation learning [14], etc. Data augmentation is also a successful tool to improve the generalization ability and robustness [6]. Wang et al. [6] demonstrate that a standard classifier can generalize well if trained with careful pre-/post-processing and data augmentation.

However, as the manipulation methods are rapidly developed and improved, how the existing detectors will generalize to various AI-synthesized images (as shown in Fig. 1) deserves deep investigation. On the one hand, advanced generation models have made the differences between real and fake images more subtle and local. On the other hand, differ-

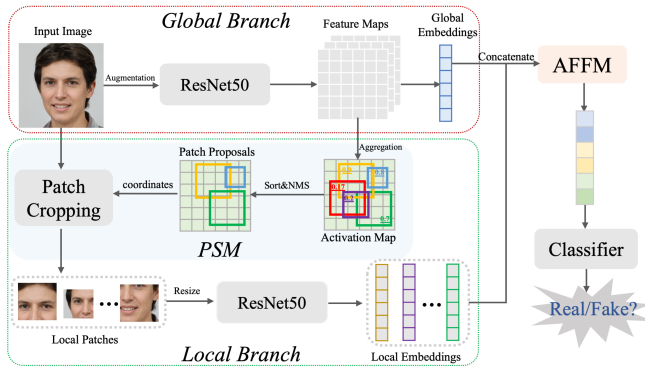


Fig. 2. The architecture of our model. Global branch extracts spatial feature and local branch extracts subtle feature from local patches selected by Patch Selection Module (PSM). The Attention-based Feature Fusion Module (AFFM) fuses global and local features for the classification.

ent types of manipulation with either global synthesis or local manipulation are increasing realistic. Thus, the methods that only make use of global spatial features may not work well on the models: 1) that synthesize entire images with refined global qualities and subtle artifacts in local regions; 2) that partially manipulate the image like swapping the face region only in recent DeepFakes.

Inspired by the previous works [7, 15], we propose a two-branch model to combine global and local features to detect diverse types of AI synthesized images. The global branch focuses on extracting overall structural features from the whole image. The local branch extracts fine-grained artifacts from multiple patches. Instead of using random or manually predefined patches, we design a Patch Selection Module (PSM) to automatically select informative patches. The features from two branches are fused by an Attention-based Feature Fusion Module (AFFM) and fed into a binary classifier.

Our main contributions can be summarized as follows:

- We propose a two-branch framework to combine global features from the whole image with the local features from informative patches to enhance the generalization ability of AI-synthesized image detection.
- We design an attention-based patch selection module to automatically select multiple patches for local subtle artifacts extraction without any manual annotations. This module facilitates the model to locate locally manipulated regions more efficiently.
- We collect a large-scale and highly diverse dataset with different resolutions, various manipulation types, and extensive synthesis models to evaluate our method. Experimental results demonstrate the outstanding generalization ability of our method.

2. RELATED WORK

Fake Image Generation. The generative adversarial networks (GANs) are usually used to synthesize entire images. Unconditional GANs take a random vector as input and no class labels are needed for generative modeling. Examples of widely used unconditional GAN models include

StyleGAN-V1/2/3 [16, 1, 17], which can generate realistic high-resolution images. Conditional GANs models such as CycleGAN [18] often take an image as input and generate the specific class of image with different object, style, etc. For partially manipulated images/videos such as face swaps, puppetry and attributes manipulation, most methods are based on the auto-encoder models [19]. Several large scale datasets such as FaceForensic++ [20] and Celeb-DF [2] are proposed.

Fake Image Detection. There have been many active developments in methods to distinguish between real and synthetic/manipulated images. A common approach is to train a binary classification detector on whole image with deep neural network models [6, 8, 9, 7, 21]. These methods focus on global feature extraction from the whole image and are often not sensitive to the subtle local artifacts. To obtain more generalized features for diverse fake manipulation, recent studies have emphasized the significance of informative patches (local features) for detection and localization [22, 23, 24]. Specifically, the study in [22] proposes a patch-based classifier with limited receptive fields and obtains excellent generalization on the entire face forgery datasets. The method in [23] also focuses on various facial parts and evaluate the performance on each individual part. However, this method puts too much emphasis on local artifacts and ignores global ones, which may not generalize well to both entire forgery and partially forgery detection.

3. METHOD

3.1. Overall Architecture

The overall structure of our two-branch model is illustrated in Fig. 2. Taking an image as input, the global branch generates a stack of global feature maps and a global feature embedding using the backbone network (e.g. ResNet50 [38]). The Patch Selection Module (PSM) predicts the coordinates of most informative patches that are then cropped from the input image and fed into local branch to obtain the local features. An Attention-based Feature Fusion Module (AFFM) is further designed to fuse the global and patch embeddings.

3.2. Global Branch

The global branch in the model is designed for high-level global spatial feature learning. It takes an image as input and employs the CNN feature extractor to obtain the global features maps from last convolutional layer, denoted as $F \in \mathbb{R}^{C \times H \times W}$ with C channels and $H \times W$ spatial size. The global embeddings are acquired by inputting F into average pooling and fully connected layers. The global features contain the spatial and structural information from the input image.

3.3. Local Branch with Patch Selection Module

Our local branch is designed to locate the most informative patches, from which the discriminative local features are extracted. Assuming that the regions with higher energy in global feature maps usually are more informative for classification task, we design a Patch Selection Module (PSM)

Table 1. Details of our highly-diverse testing dataset.

Family	Models	# Images (Real/Fake)	Object	Resolution
Contiditional GANs	GauGAN[25], StarGAN[26], CycleGAN[18], Pix2Pix[27]	19,343 (10,023/9,320)	Face & Non-face (bedroom,apple,winter,etc.)	256 × 256
Uncontiditional GANs	ProGAN[28], StyleGAN[16], StyleGAN2[1], StyleGAN3[17], BigGAN[29], WGAN[30], VAEGAN[31]	71,750 (27,335/44,415)	Face & Non-face (car,church,horse,etc.)	160 × 160 ~ 5k × 8k
Perceptual Loss	CRN[32], IMLE[33]	25,528 (12,764/12,764)	Non-face(street)	512 × 256
Low-level Vision	SITD[34], SAN[35]	798 (399/399)	Face&Non-face(flower, animal, etc.)	480 × 320 ~ 6k × 4k
DeepFakes	FF++[20], Celeb-DF[2]	7,097 (3,531/3,566)	Face	256 × 256 ~ 944 × 500
Others	Transformers[36], BGM[37]	3,908 (1,438/2,470)	Face	256 × 256 ~ 1024 × 1024

module to locate more useful patches from global feature maps. The PSM module extracts the coordinates of the most informative patches in the input image by calculating the patch scores from the global feature maps F .

To calculate the patch scores, we first estimate the activation map A by aggregating the F in the channel dimension as $A(x, y) = \sum_{j=1}^C F^j(x, y)$. Then we calculate the average score in a window on the activate map to indicate the “informativeness” of the corresponding patch in the input image. Specifically, we set the windows with $H_w \times W_w$ spatial size to slide on the activate map A and calculate the score of each window S_w as following:

$$S_w = \frac{\sum_{x=0}^{W_w-1} \sum_{y=0}^{H_w-1} A(x, y)}{W_w \times H_w} \quad (1)$$

By sorting the scores, a fixed number of windows with higher scores are selected as our patch proposals after using non-maximum suppression (NMS) to reduce patch redundancy. Coordinates of these patch proposals are then mapped back to the original input image. After patch cropping, the most informative local patches are selected, which are fed into the CNN to extract local embeddings.

3.4. Attention-based Feature Fusion Module

To efficiently combine the global and local features for final classification, an Attention-based Feature Fusion Module (AFFM) based on multi-head attention mechanism [39] is designed. After concatenating the global embeddings from global branch and several patch embeddings from local branch, the AFFM learns correlations among these embeddings and combine them as a whole feature with the learned weights for further classification. Specifically, the concatenated embeddings are inputted into AFFM as the query Q , the key K , and the value V . Multi-head attention with n heads is used. The i th attention head $h_i (i \in [1, n])$ is computed with the scaled dot product attention as $h_i = \text{softmax}(\frac{W_i^Q Q \times (W_i^K K)^T}{\sqrt{d}}) \times W_i^V V$, where W_i^Q , W_i^K , W_i^V are parameter matrices and d is the dimension of K . The attention scores of all heads are concatenated as the input of a linear transformation to obtain the value of multi-head attention as $MHA(Q, K, V) = \text{Concat}(h_i) W^o (i \in [1, n])$, where W^o is also a parameter matrix. The fused one dimensional vector from $MHA(Q, K, V)$ is taken as the learned feature and inputted into classifier for the final prediction.

4. EXPERIMENTS

4.1. Datasets

The training set provided in [6] was used for our model training. It comprises 362K real images from 20 object classes of the LSUN dataset and 362K images generated by ProGAN [28]. The spatial size of all training images is 256×256 . For the testing dataset, we composed a dataset of synthetic images generated with 19 various generation models based on several existing datasets [6, 9, 2], considering both whole image synthesis (GANs) and partial manipulations (face swaps). We also included several recently released models such as BGM [37], StyleGAN3 [17], Transformers [36], Celeb-DF [2], etc. These 19 models were further divided into 6 families based on the model structures (see more details in Table 1). A total of 128,424 images have been collected, with a high diversity in generation model, data object and image resolution. To the best of our knowledge, this is the most comprehensive and diverse dataset for AI synthesized image detection evaluation.

4.2. Implementation Details

Our model is implemented with Pytorch [40] and trained with 64 batch size and 0.0001 base learning rate. The Adam optimizer is used with the loss function of BCE. ResNet-50 [38] pre-trained with ImageNet is utilized as backbone to extract global and local features. Input image is resized to 224×224 for our global feature extraction branch. Before resizing, 10% training samples are selected randomly in each training epoch and augmented with Gaussian blur with Sigma parameter $\sigma \sim \text{Uniform}[0, 3]$ and JPEG-ed with quality $\sim \text{Uniform}\{30, 31, \dots, 100\}$. Feature maps from layer $Conv_5b$ in ResNet-50 are inputted into PSM to guide informative patches selection. In PSM implementation, two sliding windows $H_w \times W_w = 3 \times 3$ and $H_w \times W_w = 2 \times 2$ are used to capture artifacts from different resolution on activation map, and the sliding step is set as 1. 3 patch proposals from window size 3×3 are selected and mapped into 224×224 patches on the input image. Similarly, 3 patch proposals from window sizes 2×2 are mapped into 112×112 patches. These 6 patches are resized into 224 and then fed into ResNet-50 for local embeddings extraction. The dimension of extracted global embeddings and patch embeddings is 128. In AFFM, 3 attention layers with 4 heads are utilized to fuse all the extracted features.

Table 2. Evaluation results on the Seen and Unseen models. We show the mAP of each family, Total mAP of all the models and Global AP. The highest value is highlighted in black. Baselines are listed in the first three rows followed by two ablations (Our_RdmCrop and Our_Rsz448). The proposed method Our_PSM is listed in the last row.

Method \ Family	Seen		Unseen					Total	Global
	ProGAN	Conditional GANs	Unconditional GANs	Perceptual Loss	Low-level Vision	DeepFakes	Others	mAP	AP
CNN_aug[6]	100	94.782	88.100	98.374	82.027	69.576	71.846	86.389	93.578
No_down[8]	100	91.294	90.438	93.820	89.853	72.033	81.445	88.063	91.593
Patch_forensics[22]	94.620	89.245	69.286	99.472	69.661	57.424	57.902	73.979	70.149
Our_RdmCrop	100	96.039	88.775	98.719	71.773	63.154	82.357	86.825	94.457
Our_Rsz448	100	97.286	92.712	99.930	74.340	68.960	80.143	88.813	95.717
Our_PSM	100	98.305	95.485	99.147	85.104	70.109	87.493	91.732	96.906


As for baselines, we use CNN_aug [6], No_down [8] and Patch_forensics [22]. We follow the settings in their paper and train them on our training dataset from scratch. Note that for Patch_forensics, we use their released weights. For metric, we utilize average precision (AP) to evaluate our model and the baselines following recent works [6]. We report the mean AP (mAP) by averaging the AP scores over each family.

4.3. Evaluation for AI-synthesized Image Detection

Table 2 reports the mAP of each family of the baselines and the proposed Our_PSM. The total mAP among all the testing models are also listed. More importantly, to accurately measure the model performances in a real-world scenario where the source of the images is unknown, we mixed up all the images in the testing dataset and reported the global AP on the mixed dataset. All the methods achieve 100% accuracy on the seen ProGAN model except Patch_forensics, which is limited to face object detection and not works well on our testing data with diverse objects. For the unseen models, compared with three baselines, Our-PSM achieves the best mAP values among most families and performs slightly worse on Low-level Vision and DeepFakes than the No_down method. We conjecture that the resizing operations of our model on the high-resolution images in Low-level Vision and DeepFakes families will make the artifacts more difficult to detect. Overall, Our_PSM achieves the best performance in terms of Total mAP and Global AP. This demonstrates the better generalization ability of our model than the methods that only use whole image information or only focus on features from small patches.

To further evaluate the effectiveness of the proposed Patch Selection Module (PSM), we also show the mAP obtained by replacing our PSM with two popular patch selection methods in the bottom 3 lines of Table 2. In Our_RdmCrop, the global feature extraction and AFFM stay the same except that the PSM is replaced with random patch cropping on the input image. Similarly, in Our_Rsz448, we replace the PSM with resizing the input image to 448 first and cropping five 224×224 patches from four corners and center. As shown in Table 2, Our_PSM outperforms other two patch selection methods in all families except in Perceptual Loss. However, considering the overall performance, Our_PSM demonstrates the best

performance because our patch selection focuses on the most distinguishable patches and achieves the balance between using all the information of image and training efficiency.



(a) Robustness to Blurring (b) Robustness to JPEG
Fig. 3. Comparison of detection robustness to post-processing.

We also evaluate the robustness of our method and the baselines on the images post-processed with Gaussian Blur and JPEG Compression. Fig. 3 shows that our model achieves good robustness to these post-processing operations even on such a diverse unseen dataset. This indicates that data augmentation during training process may help to improve the robustness.

5. CONCLUSION

In this work, we propose a two-branch model that fuses global and informative local features for generalized AI-synthesized image detection. Global branch learns the whole image structure information, and the local branch extracts local subtle features from the informative regions selected by the proposed PSM. Multi-head attention is finally used to fuse these complementary features for a binary classification. Our model learns to focus on the most discriminative local patches with the help of the global high-level features in an unsupervised way, which allows us to fuse the local and global information more efficiently. Extensive experiments on the highly diverse testing dataset that contains 19 models demonstrate that our proposed method achieves high accuracy and good generalization ability. In future work, we will investigate a differentiable patch selection module so that the patches can be learned more efficiently.

Acknowledgements: This work is supported by the US Defense Advanced Research Projects Agency (DARPA) Semantic Forensic (SemaFor) program, under Contract No. HR001120C0123. We thank SemaFor for providing datasets for our training and testing.

6. REFERENCES

- [1] Tero Karras et al., “Analyzing and improving the image quality of stylegan,” in *CVPR*, 2020, pp. 8110–8119.
- [2] Yuezun Li et al., “Celeb-df: A large-scale challenging dataset for deepfake forensics,” in *CVPR*, 2020, pp. 3207–3216.
- [3] Hui Guo et al., “Open-eye: An open platform to study human performance on identifying ai-synthesized faces,” *arXiv preprint arXiv:2205.06680*, 2022.
- [4] Xin Wang et al., “Gan-generated faces detection: A survey and new perspectives,” *arXiv preprint arXiv:2202.07145*, 2022.
- [5] Siwei Lyu, “Deepfake detection: Current challenges and next steps,” in *2020 IEEE international conference on multimedia & expo workshops (ICMEW)*. IEEE, 2020, pp. 1–6.
- [6] Sheng-Yu Wang et al., “Cnn-generated images are surprisingly easy to spot... for now,” in *CVPR*, 2020, pp. 8695–8704.
- [7] Hanqing Zhao et al., “Multi-attentional deepfake detection,” in *CVPR*, 2021, pp. 2185–2194.
- [8] Diego Gragnaniello et al., “Are gan generated images easy to detect? a critical analysis of the state-of-the-art,” in *ICME*. IEEE, 2021, pp. 1–6.
- [9] Vishal Asnani et al., “Reverse engineering of generative models: Inferring model hyperparameters from generated images,” *arXiv preprint arXiv:2106.07873*, 2021.
- [10] Nils Hulzebosch et al., “Detecting cnn-generated facial images in real-world scenarios,” in *CVPRW*, 2020, pp. 642–643.
- [11] Hyeonseong Jeon et al., “T-gd: Transferable gan-generated images detection framework,” *arXiv preprint arXiv:2008.04115*, 2020.
- [12] Francesco Marra et al., “Incremental learning for the detection and classification of gan-generated images,” in *WIFS*. IEEE, 2019, pp. 1–6.
- [13] Davide Cozzolino et al., “Towards universal gan image detection,” in *VCIP*. IEEE, 2021, pp. 1–5.
- [14] Minha Kim et al., “Fretal: Generalizing deepfake detection using knowledge distillation and representation learning,” in *CVPR*, 2021, pp. 1001–1012.
- [15] Fan Zhang et al., “Multi-branch and multi-scale attention learning for fine-grained visual categorization,” in *ICMM*. Springer, 2021, pp. 136–147.
- [16] Tero Karras et al., “A style-based generator architecture for generative adversarial networks,” in *CVPR*, 2019, pp. 4401–4410.
- [17] Tero Karras et al., “Alias-free generative adversarial networks,” *NeurIPS*, vol. 34, 2021.
- [18] Jun-Yan Zhu et al., “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *ICCV*, 2017, pp. 2223–2232.
- [19] Yisroel Mirsky and Wenke Lee, “The creation and detection of deepfakes: A survey,” *ACM Computing Surveys (CSUR)*, vol. 54, no. 1, pp. 1–41, 2021.
- [20] Andreas Rossler et al., “Faceforensics++: Learning to detect manipulated facial images,” in *ICCV*, 2019, pp. 1–11.
- [21] Hui Guo et al., “Robust attentive deep neural network for exposing gan-generated faces,” *arXiv preprint arXiv:2109.02167*, 2021.
- [22] Lucy Chai et al., “What makes fake images detectable? understanding properties that generalize,” in *ECCV*. Springer, 2020, pp. 103–120.
- [23] Steven Schwarcz et al., “Finding facial forgery artifacts with parts-based detectors,” in *CVPR*, 2021, pp. 933–942.
- [24] Xueqi Zhang et al., “Thinking in patch: Towards generalizable forgery detection with patch transformation,” in *Pacific Rim International Conference on Artificial Intelligence*. Springer, 2021, pp. 337–352.
- [25] Taesung Park et al., “Semantic image synthesis with spatially-adaptive normalization,” in *CVPR*, 2019, pp. 2337–2346.
- [26] Yunjey Choi et al., “Stargan: Unified generative adversarial networks for multi-domain image-to-image translation,” in *CVPR*, 2018, pp. 8789–8797.
- [27] Phillip Isola et al., “Image-to-image translation with conditional adversarial networks,” in *CVPR*, 2017, pp. 1125–1134.
- [28] Tero Karras et al., “Progressive growing of gans for improved quality, stability, and variation,” *arXiv preprint arXiv:1710.10196*, 2017.
- [29] Andrew Brock et al., “Large scale gan training for high fidelity natural image synthesis,” *arXiv preprint arXiv:1809.11096*, 2018.
- [30] Martin Arjovsky et al., “Wasserstein generative adversarial networks,” in *ICML*. PMLR, 2017, pp. 214–223.
- [31] Anders Boesen Lindbo Larsen et al., “Autoencoding beyond pixels using a learned similarity metric,” in *ICML*. PMLR, 2016, pp. 1558–1566.
- [32] Qifeng Chen and Vladlen Koltun, “Photographic image synthesis with cascaded refinement networks,” in *ICCV*, 2017, pp. 1511–1520.
- [33] Ke Li et al., “Diverse image synthesis from semantic layouts via conditional imle,” in *ICCV*, 2019, pp. 4220–4229.
- [34] Chen Chen et al., “Learning to see in the dark,” in *CVPR*, 2018, pp. 3291–3300.
- [35] Tao Dai et al., “Second-order attention network for single image super-resolution,” in *CVPR*, 2019, pp. 11065–11074.
- [36] Patrick Esser et al., “Taming transformers for high-resolution image synthesis,” in *CVPR*, 2021, pp. 12873–12883.
- [37] Zhanghan Ke et al., “Modnet: Real-time trimap-free portrait matting via objective decomposition,” in *AAAI*, 2022.
- [38] Kaiming He et al., “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [39] Ashish Vaswani et al., “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [40] Adam Paszke et al., “Pytorch: An imperative style, high-performance deep learning library,” *NeurIPS*, vol. 32, 2019.