

Noise Doesn't Lie: Towards Universal Detection of Deep Inpainting

Ang Li¹, Qihong Ke^{1*}, Xingjun Ma^{2*}, Haiqin Weng³,
Zhiyuan Zong³, Feng Xue³ and Rui Zhang⁴

¹The University of Melbourne

²Deakin University

³Ant Group

⁴Tsinghua University

angl4@student.unimelb.edu.au, qihong.ke@unimelb.edu.au, daniel.ma@deakin.edu.au,
haiqin.wenghaiqin@antfin.com, david.zzy@antgroup.com, gkn1fexxx@gmail.com, rayteam@yeah.net

Abstract

Deep image inpainting aims to restore damaged or missing regions in an image with realistic contents. While having a wide range of applications such as object removal and image recovery, deep inpainting techniques also have the risk of being manipulated for image forgery. A promising countermeasure against such forgeries is deep inpainting detection, which aims to locate the inpainted regions in an image. In this paper, we make the first attempt towards *universal* detection of deep inpainting, where the detection network can generalize well when detecting *different* deep inpainting methods. To this end, we first propose a novel data generation approach to generate a *universal* training dataset, which imitates the noise discrepancies exist in real versus inpainted image contents to train universal detectors. We then design a Noise-Image Cross-fusion Network (NIX-Net) to effectively exploit the discriminative information contained in both the images and their noise patterns. We empirically show, on multiple benchmark datasets, that our approach outperforms existing detection methods by a large margin and generalize well to unseen deep inpainting techniques. Our universal training dataset can also significantly boost the generalizability of existing detection methods.

1 Introduction

Image Inpainting is the process of restoring damaged or missing regions of a given image based on the information of the undamaged regions. It has a wide range of real-world applications such as the restoration of damaged images and the removal of unwanted objects. So far, plenty of inpainting approaches have been proposed, among which generative adversarial networks (GANs) [Goodfellow *et al.*, 2014] based deep inpainting techniques [Pathak *et al.*, 2016; Iizuka *et al.*, 2017; Yu *et al.*, 2018; Li *et al.*, 2019; Yu *et al.*, 2019; Li *et al.*, 2020a] have been demonstrated to be the most effective ones. One distinguished advantage of deep inpainting models is the



(a) Original image (b) Inpainted image (c) Inpainting mask (d) Detected mask

Figure 1: (a) The original image; (b) The inpainted image by the deep inpainting method in [Yu *et al.*, 2019]; (c) The mask that defines the inpainting region; (d) The detected mask by our approach.

ability to adaptively predict semantic structures and produce super realistic and fine-detailed textures.

However, like a two-edged sword cuts both ways, deep inpainting techniques come along with the risk of being manipulated for image forgery. Due to the super realistic inpainting effects, these techniques can be easily applied to replace the critical objects in an image with fake contents, and the tampered image may appear as photo-realistic as real images. Figure 1 (b) shows one such example crafted from the real image in Figure 1 (a) by a recent deep inpainting method [Yu *et al.*, 2019]. Inpainted images can potentially be used to create fake news, spread rumors on the internet or even fabricate false evidences. It is thus imperative to develop detection algorithms to identify whether and more importantly *where* an image has been modified by deep inpainting. Specifically, the goal of deep inpainting detection is to locate the exact inpainted regions in an image, as shown in Figure 1 (d).

While Li *et al.* [2019] recently proposed the first method for deep inpainting detection, its effectiveness is restricted to the inpainting technique the detector was trained on and does not generalize well to other inpainting techniques. However, in real-world scenarios, the exact techniques used to inpaint the images are often unknown. In this paper, we aim to address this generalization limitation and introduce a *universal* detector that works well even on unseen deep inpainting techniques.

Our approach is motivated by one important yet so far overlooked common characteristic of all deep inpainting methods: the patterns of the noise exists in real and synthesized contents are different. It has been shown that conventional image acquisition devices (e.g. camera sensors) leave dis-

*Corresponding authors.




Figure 2: Noise analysis on an inpainted image. (a) The image with the inpainting mask (the white region); (b) The inpainted image by the method in [Yu *et al.*, 2019]; (c) The pixel-wise difference between (b) and a denoised version of (b), i.e., a noise residual pattern.

tinctive noise patterns on each image [Farid, 2009], however, GAN-generated contents do not have the same type of noise patterns [Marra *et al.*, 2019]. As shown in Figure 2, given an inpainted image that contains both the device-captured real pixels and GAN-generated inpainted pixels, there are obvious discrepancies between inpainted region and non-inpainted region in the noise residual (i.e., pixel-wise subtraction between the inpainted image and its denoised version by a denoising filter). Motivated by this observation, we propose a universal deep inpainting detection framework with the following two important designs. First, we present a novel way of generating universal dataset for training universal detectors. The universal training dataset contains: 1) images with synthesized (using a pre-trained autoencoder) contents of arbitrary shapes at random locations, and 2) the corresponding masks. It imitates the noise pattern discrepancies between the real and synthesized contents in a more principled manner, without using any specific deep inpainting methods. Second, we propose a Noise-Image Cross-fusion Network (NIX-Net) to effectively exploit the discriminative information contained in both the images and their noise residuals. After training on the universal dataset, our NIX-Net can reliably recognize the regions inpainted by different deep inpainting methods.

In summary, our main contributions are:

- We propose a novel framework for universal deep inpainting detection, which consists of 1) a new method of generating universal training data, and 2) a two-stream multi-scale Noise-Image Cross-fusion detection Network (NIX-Net).
- We empirically show, on multiple benchmark datasets, that our proposed approach can consistently outperform existing detection methods, especially when applied to detect unseen deep inpainting techniques.
- Our universal training dataset can also improve the generalizability of existing detection methods, making it an indispensable part of future detection methods.

2 Related Work

2.1 Deep Image Inpainting

Different from conventional image inpainting approaches, deep learning based image inpainting (or deep inpainting for short) trains inpainting networks on large-scale datasets and can generate more visually plausible details or fill large missing regions with new contents that never exist in the input image. By far, the generative adversarial networks (GAN)

[Goodfellow *et al.*, 2014] based inpainting methods are arguably the most powerful methods for deep image inpainting. These methods all employ a GAN-based training approach with two sub-networks: an inpainting network and a discriminative network. The former learns image semantics and fills the missing regions with predicted contents, whereas the latter distinguishes whether the image is real or inpainted. Phatak *et al.* [2016] proposed the Context-Encoder (CE) for single image inpainting, which is known as the first GAN-based image inpainting technique. This technique was later improved by [Iizuka *et al.*, 2017] using dilated convolution and global-local adversarial training. Yu *et al.* [2018] proposed a two-stage inpainting network with a coarse-to-fine learning strategy. A gated convolution network along with a learnable dynamic feature selection mechanism (for each channel and at each spatial location) was proposed in [Yu *et al.*, 2019] for image inpainting. Li *et al.* [2020a] devised the Recurrent Feature Reasoning network which recurrently enriches information for the hole region. Despite the diversity of existing deep inpainting methods, they all share a common characteristic: the noise patterns in generated contents are different from those in the real image contents. While this characteristic has been observed in understanding the artificial fingerprints of GANs [Marra *et al.*, 2019], it has not been exploited for deep inpainting detection. In this paper, we will leverage such a universal characteristic of generated contents to build universal deep inpainting detectors.

2.2 Inpainting Forensics

Deep inpainting detection falls into the general scope of image forensics, but quite different from the conventional image manipulation detection or deepfake detection. Conventional image manipulation detection deals with traditional image forgery operations such as splicing [Huh *et al.*, 2018] and copy-move [Wu *et al.*, 2018]. Deepfake (or deep face swapping) is the other type of deep learning forgery techniques that swaps one person’s face in a video to that of a different person, which often requires heavy post-processing including color transfer and boundary blending [Li *et al.*, 2020b]. Different from conventional image manipulation or deepfake, deep inpainting takes one image and a mask as inputs and generates new content for the mask region, based on information of the non-mask regions within the same image. In this paper, we focus on deep image inpainting detection and the generalizability of the detector to unseen deep inpainting techniques (not to conventional image manipulation or deepfake).

Most of existing inpainting forensic methods are developed to detect traditional image inpainting techniques. For example, the detection of traditional diffusion-based inpainting based on local variance of image Laplacian [Li *et al.*, 2017], and the detection of traditional patch-based inpainting via patch similarities computed by zero-connectivity length [Wu *et al.*, 2008], two-stage suspicious region search [Chang *et al.*, 2013] or CNN-based encoder-decoder detection networks [Zhu *et al.*, 2018]. These methods are generally less effective on deep inpainting techniques that can synthesize extremely photo-realistic contents or new objects that never exist in the original image. Deep inpainting detection is a

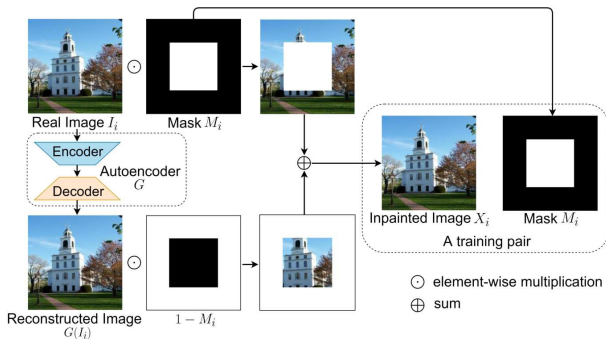


Figure 3: The procedure of universal training dataset generation.




Figure 4: The procedure of training the autoencoder G .

fairly new research topic. The high-pass fully convolutional network by [Li and Huang, 2019] is the only known detection method for deep inpainting detection. However, training this detection network requires knowing the deep inpainting technique to be detected. This tends to limit its generalizability to unseen deep inpainting methods, as shown in our experiments. In this paper, we focus on developing universal deep inpainting detectors, which we believe is a crucial step towards more powerful and practical deep inpainting detectors.

3 Problem Formulation and Universal Training Dataset Generation

3.1 Problem Formulation

Given an image X_i inpainted by a certain deep inpainting method on regions defined by a binary mask M_i , deep inpainting detection aims to locate the inpainted regions M_i . A detection network can be trained to take the inpainted image X_i as input and output the predicted mask \hat{M}_i .

To train the detector, a straightforward approach is to use inpainted images generated by a deep inpainting method as the training data [Li and Huang, 2019]. However, this *inpainting-method-aware* approach often generalizes poorly when applied to detect inpainted images generated by unseen deep inpainting methods. An empirically analysis can be found in Section 5.1. In contrast, we propose to generate a universal training dataset to capture the common characteristics shared by different deep inpainting methods, and train the detection model on this universal dataset. Such an *inpainting-method-agnostic* approach can improve the generalization of the detection model to unseen deep inpainting methods. Next, we will introduce our proposed universal training dataset gen-

eration method inspired by the distinctive noise patterns exist in real versus generated contents.

3.2 Universal Training Dataset Generation

A universal training dataset should consider the common characteristics of different deep inpainting methods, rather than relying on the specific artifacts of one particular deep inpainting method. Motivated by our observation in Figure 2, here we propose to create the universal training dataset by simulating “inpainted” images from autoencoder reconstructions, instead of using any existing deep inpainting methods. The complete generation procedure is illustrated in Figure 3.

For a set of real images $I := \{I_1, I_2, \dots, I_n\}$, the generation process generates a set of simulated (as it does not use any real inpainting methods) inpainted images $X := \{X_1, X_2, \dots, X_n\}$ using a pre-trained autoencoder with a set of random binary masks $M := \{M_1, M_2, \dots, M_n\}$. Specifically, given a real image I_i , we first obtain its reconstructed version $G(I_i)$ from an autoencoder G . We train the autoencoder following a typical GAN [Goodfellow *et al.*, 2014] approach using the autoencoder as the GAN generator and an additional classification network as the GAN discriminator. The overall structure is illustrated in Figure 4. The autoencoder is trained to have small reconstruction error, and at the same time, the reconstructed images should be as realistic as the real images according to the discriminator. The overall training loss of this autoencoder is:

$$\mathcal{L} = \sum_{i=1}^n \underbrace{\log(D(I_i)) + \log(1 - D(G(I_i)))}_{\mathcal{L}_{adv}} + \lambda \underbrace{\|G(I_i) - I_i\|_2}_{\mathcal{L}_{rec}} \quad (1)$$

where, \mathcal{L}_{adv} and \mathcal{L}_{rec} are the adversarial and reconstruction loss respectively, and $\lambda = 0.1$ is a trade-off parameter.

After training, the autoencoder G is applied to reconstruct each real image in I . With the reconstructed images, our next step is to simulate the inpainting process. Specifically, we simulate an “inpainted” image X_i by combining I_i and its reconstruction $G(I_i)$ according to a random mask M_i :

$$X_i = M_i \odot I_i + (1 - M_i) \odot G(I_i), \quad (2)$$

where \odot is the element-wise multiplication and M_i is a binary mask with 0 elements indicating the inpainting region (white region) and 1 elements indicating the non-inpainting region (black region). In X_i , the “inpainted” region carries over the noise patterns of the synthesized contents from the autoencoder, while the rest of the regions preserve the noise patterns from the real contents. Following this procedure, we can obtain a set of images with synthesized regions to create the universal training dataset: $UT = \{X, M\}$, with X being the simulated inpainted images and M being the inpainting masks.

Our universal dataset UT distinguishes itself from existing inpainting-method-aware datasets as a general formulation of real versus generated contents. However, having this dataset is not enough to train accurate deep inpainting detectors. It requires an effective learning framework to exploit the discriminative information contained in the dataset, especially the noise patterns. Next we will introduce our proposed detection network that serves this purpose.




Figure 5: Overview of the proposed Noise-Image Cross-Fusion Network (best viewed in color). Here “upsample 1×1 ” and “stride-2 3×3 ” refer to bilinear upsampling followed by a 1×1 convolution, and 3×3 convolution with stride 2, respectively.

4 Noise-Image Cross-fusion Network (NIX-Net)

As shown in Figure 2 (c), the noise residual patterns between inpainted and non-inpainted regions in an inpainted image are distinct. Unlike previous works that ignore this important cue for mask detection, we propose NIX-Net which leverages both the inpainted image X_i and its noise residual R_i to enhance detection performance. As shown in Figure 5, the proposed NIX-Net consists of three components: 1) feature extraction, 2) multi-scale cross fusion, and 3) mask detection.

Feature Extraction. Given an inpainted image X_i , we define its noise residual as following:

$$R_i = X_i - d(X_i), \quad (3)$$

Inspired by recent studies on image forensics using SRM features [Zhou *et al.*, 2018], we choose the SRM filter as our denoising filter $d(\cdot)$. The feature extraction component includes two parallel feature extraction streams to learn multi-scale feature maps of the input image and its noise residual, respectively. Each feature extraction module consists of three ResNet blocks [He *et al.*, 2016b], resulting in three feature maps of X_i , namely, $\phi_{I_1}(X_i)$, $\phi_{I_2}(X_i)$ and $\phi_{I_3}(X_i)$. The numbers of channels in the feature maps are 128, 256 and 512, respectively. Likewise, the three feature maps of R_i are $\phi_{N_1}(R_i)$, $\phi_{N_2}(R_i)$ and $\phi_{N_3}(R_i)$. Since the last convolution layer of each ResNet block has a stride of 2 to reduce the spatial scale [He *et al.*, 2016a], the spatial scales of the feature maps are $1/2$, $1/4$ and $1/8$ of the input spatial size.

Multi-Scale Cross Fusion. The multi-scale cross fusion component aims to effectively incorporate noise features and image features in a multi-scale manner. Specifically, both the noise stream and the image stream are followed by a fusion module. The fusion module takes the three-scale feature

maps as input, and outputs the crossly fused three-scale feature maps. Each fused feature map is the sum of the three transformed (upsampled, downsampled or unchanged) input feature maps. The outputs of fusion module 1 in the image stream are $\phi_{I_1}^F(X_i)$, $\phi_{I_2}^F(X_i)$ and $\phi_{I_3}^F(X_i)$, while the outputs of fusion module 2 in the noise stream are $\phi_{N_1}^F(R_i)$, $\phi_{N_2}^F(R_i)$ and $\phi_{N_3}^F(R_i)$. The purpose of this fusion module is to exchange the information across multi-scale feature representations, and produce richer feature representations with strengthened position sensitivity [Sun *et al.*, 2019]. For each scale, we concatenate the feature maps of the image stream and the noise stream along the channel dimension, followed by a Conv block. The outputs are denoted as $\psi_1(X_i, R_i)$, $\psi_2(X_i, R_i)$ and $\psi_3(X_i, R_i)$. The Conv block consists of two 3×3 convolutions, each of which is followed by a batch normalization and a ReLU Layer. Finally, fusion module 3 further consolidates the connection between the image and noise features over different scales, and output three fused feature maps $\psi_1^F(X_i, R_i)$, $\psi_2^F(X_i, R_i)$ and $\psi_3^F(X_i, R_i)$.

Mask Detection. The mask detection module first upsamples the two lower-resolution feature maps $\psi_2^F(X_i, R_i)$ and $\psi_3^F(X_i, R_i)$ using bilinear upsampling so that they have the same resolution as $\psi_1^F(X_i, R_i)$. The three feature maps are then concatenated along the channel dimension, followed by a Conv block and an upsampling layer which outputs a $w \times h \times 1$ feature map. The feature map is then fed into a Sigmoid layer for classification, rendering the possibility map P_i with pixel-wise predictions. Finally, the detected mask \hat{M}_i can be obtained by binarizing P_i according to a threshold value. In this paper, we set the threshold value as 0.5.

Network Training. We train the entire network end-to-end using the focal loss [Lin *et al.*, 2017] on the universal training

dataset *UT*. Note that the network can also be trained on any other inpainting detection datasets including the inpainting-method-aware datasets. The use of the focal loss is to mitigate the effect of class imbalance (the inpainted regions are often small compared to the entire image). The focal loss \mathcal{L}_f is defined as following:

$$\mathcal{L}_f = \sum_i^n (-M_i(1 - \hat{M}_i)^\gamma \log \hat{M}_i - (1 - M_i)\hat{M}_i^\gamma \log(1 - \hat{M}_i)) \quad (4)$$

where γ is the focusing parameter and is set to 2.

5 Experiments

In this section, we first introduce the experimental settings, then evaluate the performance of our proposed approach via extensive experiments and ablation studies.

Inpainting Methods and Datasets. We use three different deep inpainting techniques including **GL** [Iizuka *et al.*, 2017], **CA** [Yu *et al.*, 2018] and **GC** [Yu *et al.*, 2019] to generate inpainted images on two datasets Places2 [Zhou *et al.*, 2017] and CelebA [Liu *et al.*, 2015]. For each of the two datasets, we randomly select (without replacement) 50K, 10K and 10K images to create the training, validation and testing subsets respectively, following either our universal data generation or using one of the above three inpainting techniques (GL, CA and GC). We train the detection models on the training subset and test their performance on the test subset.

Mask Generation. To simulate more diverse and complex real-world scenarios, we utilize the irregular mask setting in [Yu *et al.*, 2019] with arbitrary shapes and random locations for both training and testing. Besides, object-shape masks are also adopted for visual comparison, as shown in Figure 6.

Baseline Models. We consider two baseline models: 1) **LDICN** [Li and Huang, 2019], a fully convolutional network designed for deep inpainting detection; and 2) **ManTra-Net** [Wu *et al.*, 2019], a state-of-the-art detection model for traditional image forgery such as splicing.

Performance Metric. We use the Intersection over Union (IoU) as the performance metric, and report the mean IoU (mIoU) over the entire test subset of inpainted images.

Training Setting. We train the networks using the Adam optimizer with initial learning rate 1×10^{-4} . An early stopping strategy is also adopted based on the mIoU on the validation dataset: the model with the highest validation mIoU is saved as the final model. All of our experiments were run with a Nvidia Tesla V100 GPU.

5.1 Quantitative Performance Evaluation

We have 3 detection networks (LDICN, ManTra-Net and our NIX-Net) and 2 types of training data including 1) 3 *inpainting-method-aware* training datasets created using 3 inpainting methods (GL, CA and GC), and 2) our *universal* (UT) training dataset. Here, we first train the detection networks on GL/CA/GC then test the performance on the test sets of all three datasets. Besides, we run our NIX-Net on a hybrid dataset that combines UT with one out of the 3 *inpainting-method-aware* datasets. We also train the detection networks on UT only to test the importance of universal

Model	Training Data		Test mIoU					
	GL	UT	Places2			CelebA		
			GL	CA	GC	GL	CA	GC
LDICN	✓		83.47	66.70	56.24	87.27	67.61	64.16
ManTra-Net	✓		88.76	70.18	64.60	92.53	76.22	70.98
NIX-Net	✓		91.82	80.55	77.63	93.37	84.48	81.24
NIX-Net	✓	✓	92.14	86.09	81.98	93.71	89.63	87.95
	CA	UT	GL	CA	GC	GL	CA	GC
LDICN	✓		69.53	82.48	57.73	75.85	87.04	68.49
ManTra-Net	✓		76.22	86.08	69.61	81.21	89.40	77.39
NIX-Net	✓		83.57	88.75	76.49	87.93	92.30	83.77
NIX-Net	✓	✓	90.50	89.16	83.80	92.49	92.74	88.36
	GC	UT	GL	CA	GC	GL	CA	GC
LDICN	✓		70.55	68.16	84.24	77.62	73.81	87.29
ManTra-Net	✓		80.85	74.69	84.90	83.31	81.25	88.46
NIX-Net	✓		84.77	81.03	85.38	90.57	86.44	88.97
NIX-Net	✓	✓	91.48	87.25	85.61	93.11	91.82	90.34
		UT	GL	CA	GC	GL	CA	GC
LDICN		✓	82.95	80.79	78.29	85.52	82.98	81.43
ManTra-Net		✓	88.42	83.15	80.52	89.71	86.64	85.38
NIX-Net		✓	91.33	88.46	84.71	93.06	91.59	88.20

Table 1: Quantitative Comparison on Places2 and CelebA datasets.

training data to generalizability. Note that all these experiments are run separately for Places2 and CelebA. The performance of the 3 detection networks are summarized in Table 1.

Overall Performance. As shown in Table 1, our NIX-Net outperforms existing methods by a large margin in all test scenarios, especially when transferred to detect unseen inpainting methods. When trained on the hybrid datasets, our NIX-Net achieved the best overall performance. Next, we will provide a detailed analysis of these results from two perspectives: 1) the effectiveness of different detection networks, and 2) the importance of universal training data.

Effectiveness of Different Detection Networks. For LDICN and ManTra-Net, although decent results can be obtained on the known (used for generating training data) inpainting method, their performance drops drastically on unseen inpainting methods. Such a poor generalizability indicates that both models overfit to the artifacts of a particular inpainting method and fail to consider the common characteristics of different deep inpainting techniques. By contrast, our NIX-Net demonstrates consistently better generalizability, regardless of the inpainting method used for training. This is largely due to the sufficient (multi-scale and cross fusion) exploitation of the noise information contained in real versus inpainted contents. This also indicates that noise patterns are indeed a reliable cue of detecting inpainted regions.

Importance of Universal Training Data. Revisit Table 1, we find that, whenever the UT dataset is used in conjunction with one inpainting-method-aware dataset, the generalization performance of our NIX-Net can be significantly improved. Moreover, the UT dataset alone can lead to much better generalizability of existing methods LDICN and ManTra-Net. This result verifies, from the data perspective, the importance of noise modeling for universal deep inpainting detection. More importantly, such noise modeling like our proposed universal training dataset generation is much easier than generating training data using different deep inpainting techniques,

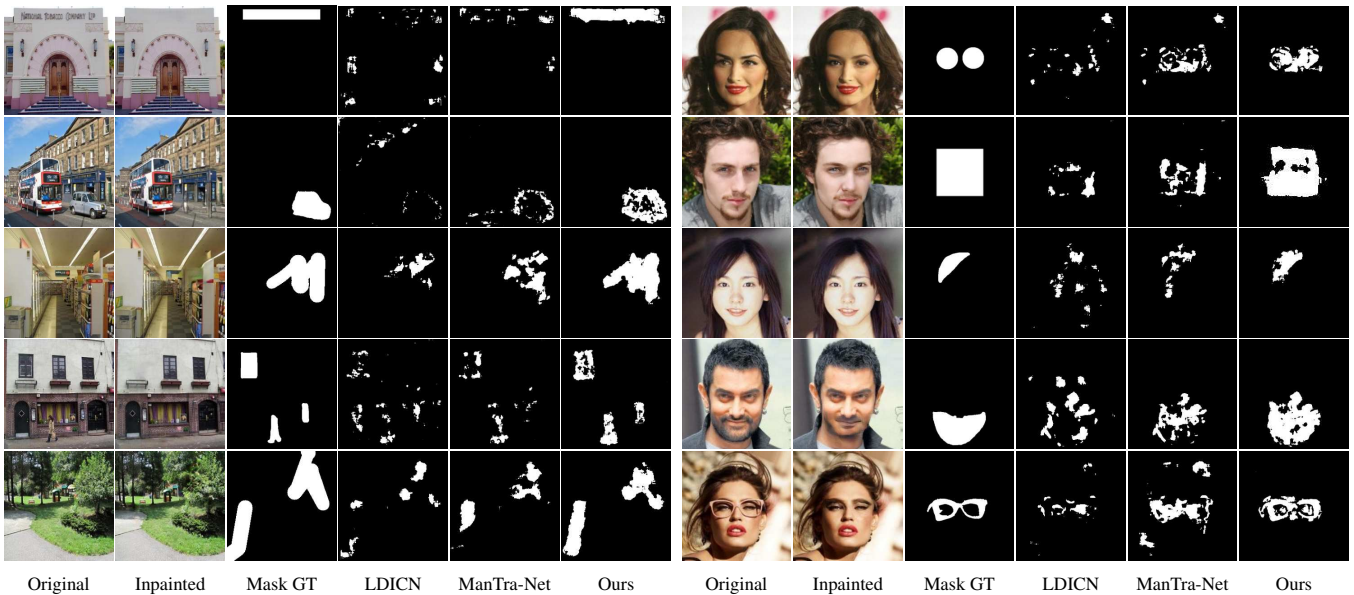


Figure 6: Qualitative comparisons on Places2 and CelebA. The original images are inpainted by CA. Mask GT refers to the ground truth of inpainting mask. LDICN and ManTra-Net are only trained on data generated by GL. Our model is only trained with UT data.

making it more practical for real-world applications. Another important observation is that NIX-Net trained on UT alone can achieve a similar level of performance as it was trained on the hybrid datasets, though combining more training data does improve the performance.

5.2 Qualitative Performance Evaluation

Here, we provide a qualitative comparison by visualizing the detected masks. Figure 6 illustrates some of the examples on Places2 and CelebA for LDICN/ManTra-Net trained only on data generated by GL and our NIX-Net trained only on UT data. The ten tested images are all inpainted by CA. By comparing the ground truth mask (Mask GT) and the masks predicted by different detection networks, one can find that our NIX-Net can produce the most similar masks to the ground truth. LDICN and ManTra-Net, however, cannot accurately identify the inpainted regions, especially when they are complex (the three bottom rows). These visual inspections confirm the superiority of our proposed universal training data generation approach and the NIX-Net detection network.

5.3 Ablation Study

Here, we run a set of ablation studies to provide a complete understanding of the two key components of our NIX-Net network: two-stream (noise+image) feature learning and multi-scale cross fusion. Table 2 compares the full NIX-Net detection network with its five variants created by removing or keeping the noise/image stream or the three fusion modules. All these networks are trained on the UT dataset generated from Places2 and tested on test images from Places2 by GL, CA and GC. It shows that, after removing either the noise or the image stream, the performance degrades drastically. This implies that both the image and the noise pattern are crucial for extracting rich features for detection. The worst

performance is observed when all 3 fusion modules removed, even though it still has the noise and the image streams. When adding either fusion module 1 and 2 or fusion module 3 back into the network, the performance is clearly improved. These results indicate that the proposed fusion module is essential for exchanging the information across multi-scale representations and achieving semantically richer feature fusion.

Ablation of NIX-Net	Test mIoU		
	GL	CA	GC
w/o noise stream	88.24	84.11	79.87
w/o image stream	83.67	77.59	74.14
w/o all fusion modules	79.36	76.22	67.84
w/o fusion module 1 and 2	84.72	82.35	77.93
w/o fusion module 3	89.19	85.35	81.44
Full NIX-Net	91.33	88.46	84.71

Table 2: Ablation of NIX-Net for (a) the noise/image stream or (2) the 3 multi-scale cross fusion modules. Networks are trained on the UT dataset generated for Places2 and tested on GL/CA/GC test sets.

6 Conclusion

In this work, we have proposed an effective approach for universal deep inpainting detection. Our approach consists of two important designs: 1) a novel universal training dataset generation method and 2) a Noise-Image Cross-fusion (NIX-Net) detection network. Extensive experiments on two benchmark datasets verify the effectiveness of our proposed approach and its superior generalization ability when applied to detect unseen deep inpainting methods. Our work not only provides a powerful universal detection method but also opens up a new direction for building more advanced universal deep inpainting detectors.

References

- [Chang *et al.*, 2013] I-Cheng Chang, J Cloud Yu, and Chih-Chuan Chang. A forgery detection algorithm for exemplar-based inpainting images using multi-region relation. *Image and Vision Computing*, 31(1):57–71, 2013.
- [Farid, 2009] Hany Farid. Image forgery detection a survey. *Signal Processing Magazine, IEEE*, 26:16 – 25, 04 2009.
- [Goodfellow *et al.*, 2014] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Proceedings of the Annual Conference on Neural Information Processing Systems*, 2014.
- [He *et al.*, 2016a] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2016.
- [He *et al.*, 2016b] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, 2016.
- [Huh *et al.*, 2018] Minyoung Huh, Andrew Liu, Andrew Owens, and Alexei A Efros. Fighting fake news: Image splice detection via learned self-consistency. In *European Conference on Computer Vision*, 2018.
- [Iizuka *et al.*, 2017] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. Globally and locally consistent image completion. *ACM Transactions on Graphics (TOG)*, 2017.
- [Li and Huang, 2019] Haodong Li and Jiwu Huang. Localization of deep inpainting using high-pass fully convolutional network. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019.
- [Li *et al.*, 2017] Haodong Li, Weiqi Luo, and Jiwu Huang. Localization of diffusion-based inpainting in digital images. *IEEE Transactions on Information Forensics and Security*, 12(12):3050–3064, 2017.
- [Li *et al.*, 2019] Ang Li, Jianzhong Qi, Rui Zhang, Xingjun Ma, and Kotagiri Ramamohanarao. Generative image inpainting with submanifold alignment. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 2019.
- [Li *et al.*, 2020a] Jingyuan Li, Ning Wang, Lefei Zhang, Bo Du, and Dacheng Tao. Recurrent feature reasoning for image inpainting. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2020.
- [Li *et al.*, 2020b] Lingzhi Li, Jianmin Bao, Ting Zhang, Hao Yang, Dong Chen, Fang Wen, and Baining Guo. Face x-ray for more general face forgery detection. In *CVPR*, 2020.
- [Lin *et al.*, 2017] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [Liu *et al.*, 2015] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015.
- [Marra *et al.*, 2019] Francesco Marra, Diego Gragnaniello, Luisa Verdoliva, and Giovanni Poggi. Do gans leave artificial fingerprints? In *IEEE Conference on Multimedia Information Processing and Retrieval*, 2019.
- [Pathak *et al.*, 2016] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2016.
- [Sun *et al.*, 2019] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2019.
- [Wu *et al.*, 2008] Qiong Wu, Shao-Jie Sun, Wei Zhu, Guo-Hui Li, and Dan Tu. Detection of digital doctoring in exemplar-based inpainted images. In *International Conference on Machine Learning and Cybernetics*, 2008.
- [Wu *et al.*, 2018] Yue Wu, Wael Abd-Almageed, and Prem Natarajan. Busternet: Detecting copy-move image forgery with source/target localization. In *European Conference on Computer Vision*, 2018.
- [Wu *et al.*, 2019] Yue Wu, Wael AbdAlmageed, and Premkumar Natarajan. Mantra-net: Manipulation tracing network for detection and localization of image forgeries with anomalous features. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2019.
- [Yu *et al.*, 2018] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018.
- [Yu *et al.*, 2019] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas Huang. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019.
- [Zhou *et al.*, 2017] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [Zhou *et al.*, 2018] Peng Zhou, Xintong Han, Vlad I Morariu, and Larry S Davis. Learning rich features for image manipulation detection. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018.
- [Zhu *et al.*, 2018] Xinshan Zhu, Yongjun Qian, X. Zhao, Biao Sun, and Ya Sun. A deep learning approach to patch-based image inpainting forensics. *Signal Process. Image Commun.*, 67:90–99, 2018.