



# SafePaint: Anti-forensic Image Inpainting with Domain Adaptation

Dunyun Chen

cs\_cdy@hnu.edu.cn

College of Computer Science and Electronic Engineering  
Hunan University, Changsha, China

Xiaoshuai Wu

shinewu@hnu.edu.cn

College of Computer Science and Electronic Engineering  
Hunan University, Changsha, China

Xin Liao\*

xinliao@hnu.edu.cn

College of Computer Science and Electronic Engineering  
Hunan University, Changsha, China

Shiwei Chen

shiwei.chen@studio.unibo.it

Department of Electrical, Electronic, and Information  
Engineering  
University of Bologna, Emilia-Romagna, Italy

## Abstract

Existing image inpainting methods have achieved remarkable accomplishments in generating visually appealing results, often accompanied by a trend toward creating more intricate structural textures. However, while these models excel at creating more realistic image content, they often leave noticeable traces of tampering, posing a significant threat to security. In this work, we take the anti-forensic capabilities into consideration, firstly proposing an end-to-end training framework for anti-forensic image inpainting named SafePaint. Specifically, we innovatively formulated image inpainting as two major tasks: semantically plausible content completion and region-wise optimization. The former is similar to current inpainting methods that aim to restore the missing regions of corrupted images. The latter, through domain adaptation, endeavors to reconcile the discrepancies between the inpainted region and the unaltered area to achieve anti-forensic goals. Through comprehensive theoretical analysis, we validate the effectiveness of domain adaptation for anti-forensic performance. Furthermore, we meticulously crafted a region-wise separated attention (RWSA) module, which not only aligns with our objective of anti-forensics but also enhances the performance of the model. Extensive qualitative and quantitative evaluations show our approach achieves comparable results to existing image inpainting methods while offering anti-forensic capabilities not available in other methods.

## CCS Concepts

• Security and privacy → Security services.

## Keywords

image inpainting, anti-forensics, domain adaptation

\*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

MM '24, October 28-November 1, 2024, Melbourne, VIC, Australia

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0686-8/24/10

<https://doi.org/10.1145/3664647.3681279>

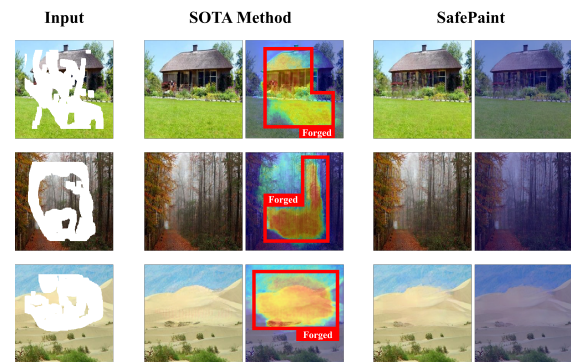


Figure 1: Examples of images inpainted by SOTA method FcF [18] and our SafePaint, which are selected from the Places2 dataset. Each result of the methods includes an inpainted image and the corresponding heatmap on inpainting detector IID-Net [34].

## ACM Reference Format:

Dunyun Chen, Xin Liao, Xiaoshuai Wu, and Shiwei Chen. 2024. SafePaint: Anti-forensic Image Inpainting with Domain Adaptation. In *Proceedings of the 32nd ACM International Conference on Multimedia (MM '24)*, October 28-November 1, 2024, Melbourne, VIC, Australia. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3664647.3681279>

## 1 Introduction

The purpose of image inpainting is to restore the missing regions of damaged images, enabling them to have semantically plausible content. Recently, image inpainting techniques have matured significantly, achieving results that are often imperceptible to the human eye. However, while inspiring, they still exhibit shortcomings. On the one hand, image inpainting techniques have positive applications, such as removing private objects from images [6]. If privacy information in images is exposed by malicious individuals through forensic methods and reverse engineering, it could be disastrous for stakeholders. On the other hand, as mentioned in [42], commonly used image evaluation metrics in this field (such as PSNR, SSIM, etc.) do not adequately represent the performance of image inpainting. An approach that performs better according to these metrics may not necessarily be visually superior to another. The ability of images

to resist forensic analysis to some extent aligns with human perceptual judgment and can serve as a basis for measuring the quality of image inpainting. Therefore, the exploration of anti-forensic image inpainting techniques is of considerable significance.

The coarse-to-fine framework [31, 36, 37, 39, 40] is a classic strategy in the field of image inpainting. It resembles human deductive reasoning, gradually restoring the image from easy to difficult, from the outer to the inner, and from superficial to deep. However, in the design of this coarse-to-fine inpainting manner, the functionalities of different networks are approximately similar, which does not fully exploit the advantages. Another research focus in image inpainting is the parallel inpainting of structure and texture [12, 18], allowing models to generate finer structural textures and make the inpainted content more realistic. However, texture structures inconsistent with the background content will attract the attention of forensic detectors, thereby posing a serious threat to the security of the inpainted image.

To further improve the performance of image inpainting methods, many approaches have incorporated attention modules into the models. Traditional spatial attention modules can enhance the network's receptive field and extraction of long-range information. However, for specific tasks such as image inpainting, their performance improvement is often limited. Yu *et al.* drew inspiration from patch-based image inpainting methods and proposed a context attention module [36, 37] tailored to the characteristics of the image inpainting tasks, simultaneously meeting the requirements of enhancing the network's receptive field and leveraging foreground information. However, its core idea of copy-move is precisely one of the key elements that forensic methods pay attention to, severely restricting the anti-forensic capabilities of the inpainting methods.

In this paper, we propose a new image inpainting architecture to achieve anti-forensic image inpainting tasks. Unlike the previous coarse-to-fine image inpainting methods that simply stack networks for multiple inpainting processes, we adopt a task-decoupled inpainting mode, using an "inpainting first, adjust later" strategy. After completing the inpainting process, we enhance the semantic alignment between the image foreground and background using the domain adaptation method, thereby improving its anti-forensic performance. Furthermore, we carefully designed a region-wise separated attention (RWSA) module, which not only fits our tasks well but also improves the performance of the model. As shown in Fig.1, compared to the state-of-the-art (SOTA) method, our approach possesses strong anti-forensic capabilities.

The main contributions are summarized as follows:

- To our best knowledge, we are the first to introduce an end-to-end training framework for anti-forensic image inpainting, making image inpainting more secure and reliable. Moreover, we innovatively incorporate the anti-forensic capabilities of image inpainting methods as one of the criteria for evaluating the quality of image inpainting, thus compensating for the shortcomings of traditional image inpainting evaluation metrics.
- We have improved the traditional coarse-to-fine image inpainting approaches by fully decoupling the image inpainting and refinement processes. This not only fully utilizes the prior knowledge obtained during the inpainting process, but

also significantly enhances the anti-forensic performance of the inpainted images through domain adaptation.

- To meet the requirements of anti-forensic image inpainting tasks and fully exploit the potential of the model, we specifically designed the region-wise separated attention (RWSA) module. This module can also be inserted into other image inpainting models to improve their anti-forensic performance.
- Comprehensive experiments on three widely-used datasets demonstrate that our proposed method significantly outperforms state-of-the-art (SOTA) methods in terms of anti-forensic capabilities, while achieving comparable performance with them in traditional evaluation metrics.

## 2 Related work

### 2.1 Traditional Inpainting Methods

**Diffusion-Based methods.** Diffusion-based methods [2, 4, 33] refer to the utilization of mathematical or physical partial differential equations (PDEs) to smoothly propagate known pixel values from the image's existing regions into the missing areas to restore damaged images.

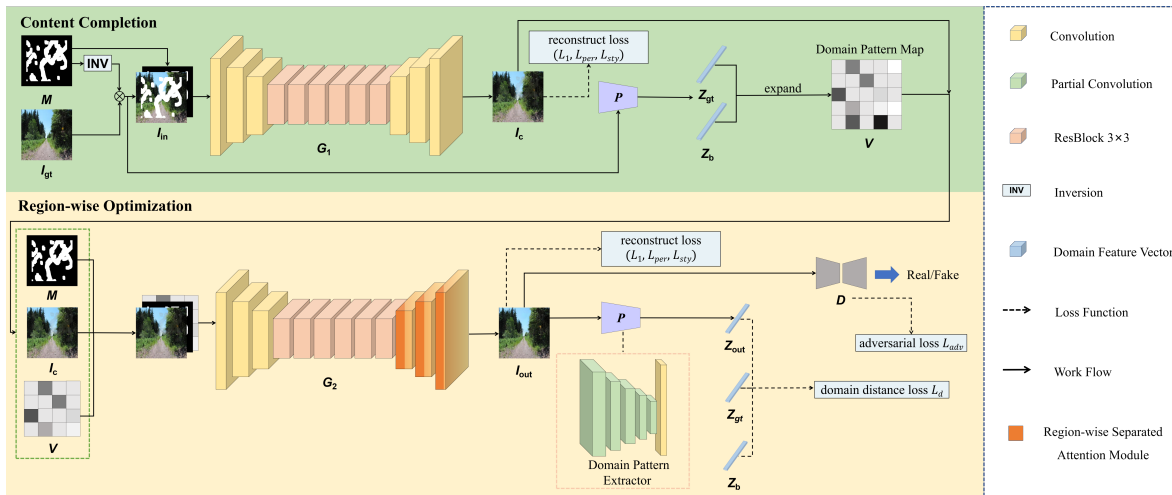
**Patch-Based methods.** Patch-based image inpainting methods [3, 5–8, 15] involve calculating and searching for samples within the known regions of the damaged images that bear the highest similarity to the missing areas, thereby utilizing them as materials to restore the damaged image. These methods exhibited strong performance in early image inpainting tasks.

However, due to significant time costs and limitations imposed by the limited information of a single image, which makes it challenging to reconstruct semantically rich images, these approaches have gradually been abandoned.

### 2.2 Deep-Learning-Based Inpainting Methods

From the standpoint of design strategy, existing deep-learning-based image inpainting methods can be roughly divided into two types: one-stage and coarse-to-fine methods.

**One-Stage Methods.** Pathak *et al.* [30] were the first to apply the generative adversarial network (GAN) [10] to image inpainting, laying the foundation for using deep learning to tackle image inpainting problems. Iizuka *et al.* [16] improved the consistency of image completion by incorporating global and local discriminators into the network. To address the limitation of traditional convolutional methods, Liu *et al.* [23] proposed partial convolution, which dynamically updates the mask during convolution computation, achieving favorable results for corrupted images with irregularly shaped masks. Zeng *et al.* [38] enhances the contextual reasoning capabilities of the network through carefully designed AOT blocks. To generate content where structure and texture align more closely, Guo *et al.* [12] adequately models both types of information in a coupled manner. Inpainting high-resolution images with large missing areas was a challenge. Therefore, Suvorov *et al.* [32] introduced a network structure based on Fast Fourier convolution (FFC), which has a larger receptive field in the initial stage of the network and performs well in complex image inpainting tasks. Inspired by this work, Jain *et al.* [18] combined the receptive power of Fast Fourier convolution with a co-modulated generator, addressing image inpainting problems by simultaneously considering image structures



**Figure 2: The overview of our proposed SafePaint. SafePaint adopts an end-to-end architecture that involves two phases. The first stage mainly focuses on content completion, while the second stage is responsible for region-wise optimization, which is the key to enhancing anti-forensic abilities.**

and textures. Due to insufficient prior knowledge, these methods are sometimes plagued by visible artifacts.

**Coarse-to-Fine Methods.** Attention mechanism enables image inpainting models to leverage long-range feature information. Inspired by patch-based methods, Yu *et al.* [36] proposed contextual attention. This module divides the image into several pixel patches and calculates their cosine similarity to guide the inpainting process. Building upon [36], Yu *et al.* [37] introduced gated convolution to address the shortcomings of partial convolution from [23]. Simultaneously, the authors introduced a sample discriminator SN-PatchGAN to generate high-quality inpainting results and accelerate training. Liu *et al.* [24] extended the attention mechanism to missing pixels. Progressive inpainting strategies have become another major technological innovation in the field of image inpainting. Zhang *et al.* [41] divided the image inpainting process into four phases and used a long short-term memory (LSTM) architecture [13] to control the information flow of the progressive process. Nazari *et al.* [29] drew inspiration from sketching, first generating a hypothetical edge image of the image to be inpainted as prior knowledge, and then filling in the missing region in the next step. Observing that there still exists useful information in the failure cases, Zeng *et al.* [40] proposed an iterative method with feedback mechanisms for inpainting. Reference [45] improved upon the limitations of the specialized convolution in [23, 37], allowing for a more flexible implementation of coarse-to-fine inpainting. Lugmayr *et al.* [27] introduced diffusion models into the field of inpainting, achieving remarkable success. Although these methods provide valuable insights into image inpainting, they fail to meet the requirements of anti-forensic image inpainting tasks.

### 2.3 Inpainting Detection Methods

While image inpainting methods are rapidly developing, the advancement of their counterpart, image inpainting detection methods, is also noteworthy. Conventional methods [22, 35] mostly focus

on manual features, *e.g.*, the similarities between image patches. Such methods incur high computational cost and often yield minimal effectiveness when confronted with complex scenarios. Recently, Many methods [21, 34] have shifted their focus toward deep-learning-based inpainting techniques. By automatic feature extraction, these methods are capable of achieving precise detection in even more complex situations. Moreover, many methods [11, 25] are proposed to detect complicated combinations of forgeries (including inpainting).

These detection methods start from a fundamentally similar point, which is leveraging the inconsistency between the distribution of generated data and that of original data.

### 2.4 Domain Adaptation

The objective of domain adaptation is to achieve the transformation from a source domain to a target domain. The pioneering work Pix2Pix [17] demonstrated remarkable results in image-to-image (I2I) translation tasks by leveraging paired images and conditional generative adversarial networks (GAN). To overcome the reliance on paired image data, Zhu *et al.* introduced CycleGAN [44]. The aforementioned methods require domain labels; in contrast, the methods [1, 19] necessitate no domain labels. In image inpainting, domain labels are typically not explicitly available. Therefore, it is proposed to use unlabeled domain adaptation aimed at enhancing the anti-forensic capabilities of our method.

## 3 Method

### 3.1 Model Architecture

As shown in Fig.2, our SafePaint achieves anti-forensic inpainting through two subtasks, *i.e.*, content completion and region-wise optimization. We will introduce the generators, the discriminator, and the domain pattern extractor of the model below.

**Generator and Discriminator.** We propose an image inpainting network that consists of two stages. It includes a generator

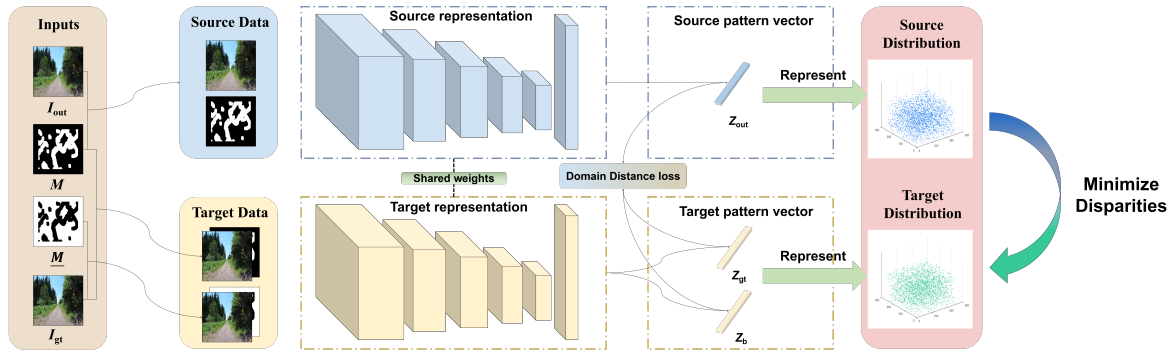


Figure 3: The implementation detail of our domain adaptation.

in each stage and a discriminator in stage *II*. The two generators  $G_1, G_2$  follow an architecture similar to the method proposed by Johnson *et al.* [20] and we use only the image generator as our backbone generator. In the upsampling process of  $G_2$ , we further integrated the RWSA module to facilitate domain adaptation and enhance performance. For the discriminator, We apply the PatchGAN [17, 44] instead of vanilla GAN to classify each image patch as pristine or inpainted. To improve the stability of the training, we also introduced Spectral normalization (SN) [28] to both the generators and the discriminator.

**Domain Pattern Extractor.** We design the domain pattern extractor to implement domain adaptation. Considering that the background and foreground regions have irregular shapes, we replaced vanilla convolutional layers with partial convolutional layers [23], which are specially designed for image inpainting with irregular masks. The output of the network will be represented in vector form, serving to characterize features from different domains.

### 3.2 Domain Adaptation for Image Inpainting

To achieve the goal of anti-forensics, we investigate the commonalities of inpainting methods and the shared properties among their corresponding forensic methods by analyzing massive corrupted image samples. It is drawn that detectors are more inclined to perceive regions where data distribution significantly deviates from that of real image data as tampered areas, a point we will further explore in Section 3.3. Hence, we employ domain adaptation as a means to narrow the gap between the foreground and background of the image, thereby evading the detection of forensic detectors.

We refer to the area to be inpainted as the foreground, and the original area as the background. Specifically, we design a domain pattern extractor to characterize the features within the delineated regions. For image inpainting, the background region is generally not allowed to be modified. In that case, the background region in images to be inpainted is considered as unified across different stages.

As shown in Fig.3, we extracted three different domain pattern vectors  $Z_b, Z_{gt}$ , and  $Z_{out}$  by the domain pattern extractor, representing background domain features, foreground domain features of ground-truth, and foreground domain features of the output image, respectively. Since domain pattern vectors can represent the data distributions within specific regions of an image, intuitively,

minimizing the disparities between them can be equivalent to reducing the distance between their corresponding domain pattern vectors. To achieve this, we design a special loss function, which will be explained specifically in Section 3.5.

### 3.3 Interpretation for Effectiveness of Domain Adaptation in Anti-forensics

In this section, we will explain how our design of domain adaptation can effectively achieve anti-forensics by analyzing the principles of inpainting methods and their respective forensic approaches. We categorize inpainting methods into diffusion-based, patch-based, and deep-learning-based methods, and we will discuss each of them separately below.

**Diffusion-Based Methods.** Taking the work [4] as an example, it leverages partial differential equations (PDEs) for image inpainting. The core idea is to gradually propagate pixels from the known regions of the image to the areas requiring inpainting. The algorithm can be formulated as follows:

$$I^{n+1}(i, j) = I^n(i, j) + \Delta v I_t^n(i, j), \forall (i, j) \in \Omega \quad (1)$$

where  $n$  denotes the number of iterations for the pixel  $I(i, j)$  at the coordinate  $(i, j)$ ,  $\Delta v$  represents the update rates,  $\Omega$  signifies the region to be inpainted, and  $I_t^n(i, j)$  is the updated value, calculated by the following formula:

$$I_t^n(i, j) = \overrightarrow{\Delta L}^n(i, j) \cdot \overrightarrow{N}^n(i, j) \quad (2)$$

where  $\overrightarrow{\Delta L}^n(i, j)$  measure the change of  $L^n(i, j)$ , which denotes the propagated information.  $\overrightarrow{N}^n(i, j)$  is the propagation direction. Since the diffusion process runs only inside  $\Omega$ , as the propagation progresses,  $I_t^n(i, j)$  will decrease gradually and converge to 0. That means  $I(i, j)$  will remain constant in the direction of  $\overrightarrow{N}$ . However, this phenomenon is extremely rare in the known regions of the image. The forensic method [22] specific to this inpainting mode confirms this, as it locates the inpainting area by comparing the local variances of the inpainted region and the known region.

**Patch-Based Methods.** The design of the work [6] is to find the most suitable patch within the known region to fill the area to be inpainted. It repeats the following three steps: (1) Computing patch priorities. (2) Propagating texture and structure information. (3) Filling region to be inpainted. We define  $\Psi_f$  as the patch with

the highest priority, and then we search in the known region for that patch  $\Psi_b$  which is most similar to  $\Psi_f$ , calculated by following rule:

$$\Psi_b = \arg \min_{\Psi_b \in \Phi} d(\Psi_f, \Psi_b) \quad (3)$$

where  $\Omega$  represents the region to be filled, the distance  $d(\Psi_m, \Psi_n)$  between two patches  $\Psi_m$  and  $\Psi_n$  is simply computed using the Euclidean distance. This rigid computational approach may introduce abnormal similarity between the inpainting region and the rest of the image, as the patch block set of the inpainting region may exhibit a one-to-many or many-to-many relationship with that of the known region. Wu *et al.* [35] utilizes this characteristic by identifying sets of block pairs exhibiting abnormal similarity as a basis for locating the area that has been inpainted.

**Deep-Learning-Based Methods.** While deep-learning-based inpainting methods mainly rely on data-driven approaches, their basic principles predominantly derive from traditional inpainting methods. For example, the contextual attention mechanism used in [36, 37] draws inspiration from patch-based inpainting methods, while the search and generate process in [24] embodies the properties of diffusion-based inpainting methods. Although these inpainting methods differ, their core ideas are nearly identical, namely, making full use of existing pixel information in images. On the other side of the coin, their corresponding forensic methods also share similarities, primarily utilizing inconsistencies of the statistical distributions between real data and generated data.

Based on the analysis above, we eventually validate the effectiveness of the proposed domain adaptation for image inpainting. Let  $X_\Omega$  the set of pixel values in the inpainting area, and  $X_\Phi$  denote the set of pixel values in the known area. Then we use  $P(X_\Omega)$  and  $P(X_\Phi)$  to represent their respective probability distribution. The Kullback-Leibler divergence (KL divergence) is utilized to express the difference between them, formulated as:

$$D_{KL}(P(X_\Phi)||P(X_\Omega)) = \int_{-\infty}^{\infty} P(X_\Phi(x)) \log \left( \frac{P(X_\Phi(x))}{P(X_\Omega(x))} \right) dx \quad (4)$$

We denote the domain adaptation transformation as function  $T$ . Our objective is to find a  $T$  that minimizes the KL divergence between these two distributions, thereby achieving the purpose of anti-forensics.

### 3.4 Region-wise Separated Attention Module

As illustrated in Fig.4, given an intermediate feature map  $x \in \mathbb{R}^{C \times H \times W}$ , and the binary mask  $M \in \mathbb{R}^{1 \times H \times W}$  of the inpainted region as the inputs (C is the feature channel, H and W is the height and width of the feature map respectively), three 1D channel attention modules  $G_b, G_{f_1}, G_{f_2}$  are utilized to weight the input features for different purposes, with the help of foreground region mask  $M$  and background region mask  $M'$ . In detail,  $G_b$  aims to select the features that are necessary for image rebuild in the background region.  $G_{f_1}$  intends to weigh the input features that have differences between the input and target. Moreover, we design the third channel attention module  $G_{f_2}$  to weight features that are unnecessary to change in the foreground region. Additionally, for the output features of  $G_{f_1}$ , we add a learnable block  $LB$  to learn the domain

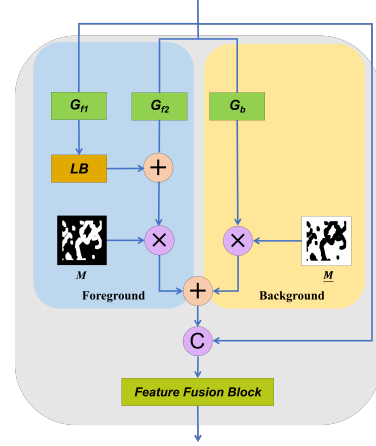


Figure 4: The detailed structure of our RWSA module.

adaptation. Therefore, the output features of foreground  $\mathcal{F}_f$  and background  $\mathcal{F}_b$  can be respectively formalized as:

$$\mathcal{F}_f = M \times [LB(G_{f_1}(x)) + G_{f_2}(x)]$$

$$\mathcal{F}_b = (1 - M) \times G_b(x) \quad (5)$$

The core component of proposed RWSA module are channel attention modules ( $G_b, G_{f_1}, G_{f_2}$ ), learnable block  $LB$ , and Feature Fusion Block  $FFB$ . Regarding the three channel attention modules, we utilize a modified Squeeze-and-Excitation Network [14] in which we add a MaxPooling layer. The learnable block  $LB$  consists of  $3 \times 3$  convolutional layers with two CONV-BN-ELU blocks to learn the textural and structural changes in the inpainted region. In order to fully integrate the feature obtained, we have designed the feature fusion block  $FFB$ , which concatenates the sum of  $\mathcal{F}_f$  and  $\mathcal{F}_b$  with the input  $x$  in channel axis. Essentially, the RWSA is to perform weighted calculations on features from different regions separately, harmonizing their differences to achieve anti-forensic goal. Above all, our region-wise separated attention (RWSA) module between input features  $x$  and output features  $\mathcal{F}$  can be finally formalized as:

$$\mathcal{F} = FFB(Concat(\mathcal{F}_f + \mathcal{F}_b, x)) \quad (6)$$

where  $Concat$  is the implementation of concatenation.

### 3.5 Loss Functions

The model training contains the  $\ell_1$  loss, perceptual loss, style loss, adversarial loss, and domain distance loss, to achieve visual plausibility and anti-forensic performance.

Formally, we designate the two generators as  $G_1$  and  $G_2$  respectively and denote the discriminator in stage II as  $D$ . Let  $I_{gt}$  be the ground truth image,  $M$  be the input binary mask which represents the region to be inpainted,  $I_{in} = I_{gt} \odot (1 - M) + M$  be the input image,  $V$  be the domain pattern map. Then the coarse result  $I_c$  and final output  $I_{out}$  are formulated as:

$$I_c = I_{gt} \odot (1 - M) + G_1(Concat(I_{in}, M)) \odot M$$

$$I_{out} = I_{gt} \odot (1 - M) + G_2(Concat(I_c, M, V)) \odot M \quad (7)$$

$\ell_1$  Loss. We adopt the  $\ell_1$  loss to calculate difference between  $I_c$  and  $I_{gt}$  as well as the difference from  $I_{out}$  to  $I_{gt}$ , formulated as:

$$\mathcal{L}_1 = \mathbb{E} [\|I_c - I_{gt}\|_1 + \|I_{out} - I_{gt}\|_1] \quad (8)$$

**Perceptual Loss.** For image inpainting, the perceptual loss from a pre-trained and fixed VGG-19 is used. The perceptual loss compares the difference between the deep feature map of the generated image and the ground truth. Such a loss function can compensate for the weakness of  $\ell_1$  Loss that helps the model capture high-level semantics. The perceptual loss is defined as:

$$\mathcal{L}_{per} = \mathbb{E} \left[ \sum_i \|\phi_i(I_c) - \phi_i(I_{gt})\|_1 + \|\phi_i(I_{out}) - \phi_i(I_{gt})\|_1 \right] \quad (9)$$

**Style Loss.** For the purpose of guaranteeing style coherence, we use style loss to compute the distance between feature maps. The style loss can be illustrated as:

$$\mathcal{L}_{sty} = \mathbb{E} \left[ \sum_i \|\psi_i(I_c) - \psi_i(I_{gt})\|_1 + \|\psi_i(I_{out}) - \psi_i(I_{gt})\|_1 \right] \quad (10)$$

where  $\psi_i(\cdot) = \phi_i(\cdot)^T \phi_i(\cdot)$  denotes the Gram matrix constructed from the activation map  $\phi_i$ .

**Adversarial Loss.** Adversarial Loss is designed to help the inpainting network  $G = \{G_1, G_2\}$  generate plausible images. Given the characteristic of our task, the discriminator  $D$  takes  $I_{gt}$  as real images and  $I_{out}$  as fake images. The adversarial loss is as follows:

$$\mathcal{L}_{adv} = \min_G \max_D \mathbb{E}_{I_{gt}} [\log D(I_{gt})] + \mathbb{E}_{I_{out}} \log [1 - D(I_{out})] \quad (11)$$

**Domain Distance Loss.** Enhancing the alignment between the foreground and background is a crucial step in our SafePaint, and we accomplish this goal by incorporating the domain distance loss. Firstly, the character  $P$  is used to denote the domain pattern extractor. As mentioned in section 3.2, the three different domain pattern vectors  $Z_b, Z_{gt}, Z_{out}$  can be computed by:

$$Z_b = P(I_{gt}, \bar{M}), \quad Z_{gt} = P(I_{gt}, M), \quad Z_{out} = P(I_{out}, M) \quad (12)$$

where  $M$  represents the input binary mask and  $\bar{M}$  represents its reverse form.

After inpainting, we attempt to make the background of the result as close to the foreground as possible while ensuring the quality of inpainted images. Therefore, we bring  $Z_{out}$  and  $Z_b$  closer while draw  $Z_{out}$  and  $Z_{gt}$  nearer by following domain distance loss:

$$\mathcal{L}_d = \mathbb{E} [d(Z_{out}, Z_b) + d(Z_{out}, Z_{gt})] \quad (13)$$

in which  $d(\cdot, \cdot)$  is Euclidean distance.

To summarize, the total loss can be written correspondingly as:

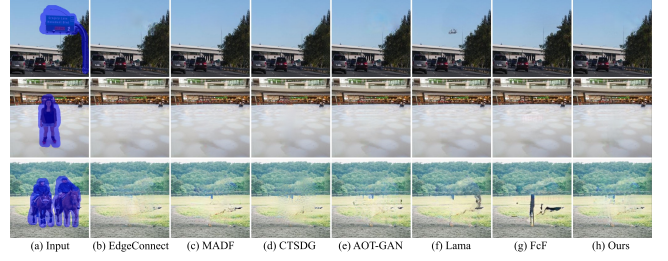
$$\mathcal{L}_{total} = \lambda_1 \mathcal{L}_1 + \lambda_{per} \mathcal{L}_{per} + \lambda_{sty} \mathcal{L}_{sty} + \lambda_{adv} \mathcal{L}_{adv} + \lambda_d \mathcal{L}_d \quad (14)$$

where  $\lambda_1, \lambda_{per}, \lambda_{sty}, \lambda_{adv}, \lambda_d$  are the weights for respective loss terms.

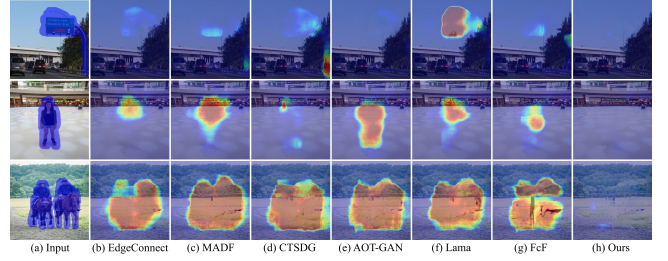
## 4 Experiments

### 4.1 Experimental Settings

We evaluate the proposed method on the Places2 [43], Paris-Street-View [9], CelebA [26] datasets, which are commonly used in the field. NVIDIA Irregular Mask [18] is used for training and testing and we classified the mask data based on their mask ratio relative



**Figure 5: Visual comparison of our SafePaint with EdgeConnect [29], MADF [45], CTSDG [12], AOT-GAN [38], Lama [32], FcF [18] on dataset Places2.**



**Figure 6: Visualization of anti-forensic capabilities comparisons on detector PSCC-Net [25].**

to the image with an increment of 10%. All the images and corresponding masks are resized to  $256 \times 256$ . We train our model in pytorch. Training is launched on a single NVIDIA RTX 4090 GPU with the batch size of 4, optimized by the Adam optimizer with  $\beta_{tas} = (0.5, 0.999)$ . We adjust the learning rate as  $2 \times 10^{-4}$  for both discriminator and generator. Besides,  $\lambda_1, \lambda_{per}, \lambda_{sty}, \lambda_{adv}, \lambda_d$  in Eq.(14) are set to 1, 0.1, 250, 0.1, 0.01. For more information of the experimental details and results, please see Appendix ?? . Code will be available at <https://github.com/hnucdy/SafePaint>.

### 4.2 Comparisons with State-of-the-art Methods

We compare our SafePaint on the Places2 [43], Paris-Street-View [9], CelebA [26] datasets with the state-of-the-art methods introduced previously: EdgeConnect [29], MADF [45], CTSDG [12], AOT-GAN [38], Lama [32], FcF [18]. We use their official implementation and the models provided by the authors. As traditional image quality assessment metrics, LPIPS and PSNR are used for quantitative quality evaluation. Since this is a piece focusing on anti-forensics, we take the anti-forensic ability of inpainting methods into account. Therefore, we adopt two pre-trained image forgery detectors PSCC-Net [25] and TruFor [11], and one inpainting-specific detector [34] for anti-forensics analysis. To quantify the anti-forensic performance, we use pixel-level Area Under Curve (AUC), F1 score on manipulation masks, and the accuracy of detectors as main metrics. Since the inpainted images serve as negative samples, lower AUC, F1, and ACC indicate better anti-forensic performance.

**Quantitative evaluation on visual quality.** Table 1 reports the numerical comparisons of proposed SafePaint with EdgeConnect [29], MADF [45], CTSDG [12], AOT-GAN [38], Lama [32], and

**Table 1: Quantitative comparison on visual quality.**

Dataset		Places2			Paris Street View			CelebA		
Mask Ratio		10%-20%	30%-40%	50%-60%	10%-20%	30%-40%	50%-60%	10%-20%	30%-40%	50%-60%
PSNR↑	EdgeConnect [29]	27.31	22.37	18.58	30.37	25.68	21.62	30.35	25.40	20.67
	MADF [45]	<b>29.18</b>	23.46	19.15	<b>32.71</b>	<b>27.15</b>	<b>22.37</b>	30.86	25.58	20.34
	CTSDG [12]	29.02	23.34	19.14	32.65	26.88	22.26	30.90	25.84	21.08
	AOT-GAN [38]	28.97	23.50	18.89	/	/	/	31.31	25.42	20.09
	Lama [32]	29.17	<b>23.56</b>	<b>19.28</b>	/	/	/	<b>32.37</b>	26.95	<b>22.50</b>
	FcF [18]	28.63	22.71	18.19	/	/	/	32.30	<b>27.01</b>	22.48
	SafePaint	29.00	23.37	19.03	32.45	27.12	22.20	32.06	26.76	21.63
LPIPS↓	EdgeConnect [29]	0.053	0.136	0.262	0.053	0.117	0.233	0.043	0.091	0.178
	MADF [45]	<b>0.037</b>	0.105	0.230	<b>0.033</b>	<b>0.090</b>	0.205	0.052	0.102	0.212
	CTSDG [12]	0.049	0.145	0.298	0.037	0.111	0.247	0.057	0.112	0.206
	AOT-GAN [38]	0.045	0.110	0.234	/	/	/	0.026	0.075	0.188
	Lama [32]	<b>0.037</b>	<b>0.102</b>	<b>0.218</b>	/	/	/	<b>0.020</b>	<b>0.054</b>	<b>0.113</b>
	FcF [18]	0.039	0.109	0.231	/	/	/	0.027	0.058	0.114
	SafePaint	0.041	0.116	0.246	0.038	0.092	<b>0.200</b>	0.032	0.068	0.138

**Table 2: Anti-forensic performance evaluation of our SafePaint with SOTAs on detector PSCC-Net [25].**

Metrics	AUC↓			F1↓			ACC↓		
Mask Ratio	10%-20%	30%-40%	50%-60%	10%-20%	30%-40%	50%-60%	10%-20%	30%-40%	50%-60%
EdgeConnect [29]	0.7376	0.7790	0.8388	0.1056	0.1437	0.2458	0.2038	0.3890	0.4970
MADF [45]	0.5738	0.8407	0.8945	0.0664	0.4724	0.7141	0.2272	0.7240	0.8988
CTSDG [12]	0.8034	0.8412	0.8533	0.2531	0.4080	0.4930	0.3060	0.6836	0.7428
AOT-GAN [38]	0.7607	0.8067	0.8661	0.1784	0.2487	0.4440	0.2724	0.5528	0.8760
Lama [32]	0.6800	0.7529	0.8084	0.0592	0.1226	0.2236	0.0790	0.2548	0.3432
FcF [18]	0.6490	0.7448	0.8084	0.0344	0.1182	0.3005	0.0598	0.5196	0.9010
SafePaint	<b>0.5628</b>	<b>0.5729</b>	<b>0.6058</b>	<b>0.0153</b>	<b>0.0092</b>	<b>0.0106</b>	<b>0.0166</b>	<b>0.0110</b>	<b>0.0132</b>

FcF [18] across the three datasets. The comparisons are based on two image quality metrics: LPIPS and PSNR. It can be observed that although we compromise visual quality to a certain extent to enhance image security in stage II, the proposed network is comparable to state-of-the-art methods in visual comparison evaluation, and the inpainted images remain visually appealing.

**Qualitative evaluation on visual quality.** Fig.5 shows the visual quality results of the SafePaint compared with existing methods on the dataset Places2. It can be seen, as one of the representative works of parallel structural and textural image inpainting, FcF [18] exhibits powerful detail generation capabilities, which sometimes paradoxically lead to content generation that is discordant with the surrounding areas. For other methods, the inpainted content may deviate from the known regions to some extent. Owing to our domain adaptation design, our SafePaint can generate more harmonious content.

**Anti-forensics analysis.** Anti-forensic performance signifies the ability of image inpainting methods to evade the detection of forensic detectors. We select two representative image forgery detectors and one inpainting detector to evaluate the anti-forensic ability of our SafePaint and the comparison methods, including PSCC-Net [25], TruFor [11] and IID-Net [34]. Table 2, Table 3 and Table 4 illustrate the comparisons regarding their anti-forensic performance on Places2. The AUC, F1, and ACC with smaller values

indicate better anti-forensic ability. We can see that our model always outperforms the state-of-the-art counterparts by a significant margin. The visual results are depicted via the heatmap in Fig.6. It is evident that compared to the other methods, our approach leaves very faint tampering traces in the image, showing excellent anti-forensic capabilities.

### 4.3 Ablation Study

In this section, we will investigate the effectiveness of each component in the SafePaint, including domain distance loss (DDL) and region-wise separated attention (RWSA) module proposed by us. Given their design for anti-forensic purposes, we will only compare the detection resistance performance of each model. Specifically, the ablation experiments are conducted on the dataset Places2 and the mask ratio is set to 30%-40%. The results of ablating each component are reported in Table 5, representing their detection evasion performance on detectors PSCC-Net, TruFor and IID-Net. We denote our SafePaint without DDL and RWSA module as "Backbone". It can be observed that "Backbone+DDL" outperforms "Backbone", demonstrating the benefit of using domain distance loss. Another observation is that "Backbone+RWSA" performs more favorably than "Backbone", which indicates the advantage of the region-wise separated attention (RWSA) module. Finally, our full method SafePaint, *i.e.*, "Backbone+DDL+RWSA", achieves the best results.

**Table 3: Anti-forensic performance evaluation of our SafePaint with SOTAs on detector Trufor [11].**

Metrics	AUC↓			F1↓			ACC↓		
	10%-20%	30%-40%	50%-60%	10%-20%	30%-40%	50%-60%	10%-20%	30%-40%	50%-60%
EdgeConnect [29]	0.8737	0.8275	0.8221	0.4432	0.3627	0.2208	0.6770	0.4304	0.2088
MADF [45]	0.8557	0.8151	0.8185	0.4175	0.3720	0.3226	0.6972	0.5200	0.3736
CTSDG [12]	0.8599	0.8249	0.8229	0.4351	0.4211	0.3340	0.7412	0.5846	0.3244
AOT-GAN [38]	0.8650	0.8232	0.8294	0.4327	0.3923	0.3012	0.7456	0.5622	0.3294
Lama [32]	0.8610	0.8216	0.8195	0.4303	0.4007	0.3528	0.7402	0.6238	0.4866
FcF [18]	0.8577	0.8172	0.8101	0.4078	0.3400	0.2314	0.6656	0.4734	0.2766
SafePaint	<b>0.6417</b>	<b>0.5781</b>	<b>0.5771</b>	<b>0.0983</b>	<b>0.0624</b>	<b>0.0620</b>	<b>0.1934</b>	<b>0.1290</b>	<b>0.1488</b>

**Table 4: Anti-forensic performance evaluation of our SafePaint with SOTAs on detector IID-Net [34].**

Metrics	AUC↓			F1↓			ACC↓		
	10%-20%	30%-40%	50%-60%	10%-20%	30%-40%	50%-60%	10%-20%	30%-40%	50%-60%
EdgeConnect [29]	0.7506	0.7845	0.7926	0.1143	0.1582	0.1743	0.3701	0.5090	0.4928
MADF [45]	0.7014	0.7234	0.7465	0.0735	0.0837	0.1149	0.1948	0.2875	0.3652
CTSDG [12]	0.7437	0.8059	0.8288	0.0921	0.2433	0.2516	0.3354	0.5978	0.7315
AOT-GAN [38]	0.8342	0.8617	0.8670	0.2709	0.3156	0.3218	0.7560	0.7913	0.8081
Lama [32]	0.7655	0.7801	0.7954	0.0836	0.1078	0.1257	0.3652	0.5484	0.6589
FcF [18]	0.7928	0.8254	0.8319	0.1624	0.2011	0.2092	0.6136	0.7032	0.7423
SafePaint	<b>0.5909</b>	<b>0.6181</b>	<b>0.6195</b>	<b>0.0398</b>	<b>0.0699</b>	<b>0.0722</b>	<b>0.0255</b>	<b>0.0283</b>	<b>0.0319</b>

**Table 5: Ablation study conducted on three detectors.**

Detectors	Model	AUC↓	F1↓	ACC↓
PSCC-Net	Backbone	0.5810	0.0395	0.5054
	Backbone+DDL	0.5753	0.0119	0.0137
	Backbone+RWSA	0.5778	0.0153	0.0182
	Backbone+DDL+RWSA	<b>0.5729</b>	<b>0.0092</b>	<b>0.0110</b>
Trufor	Backbone	0.6223	0.0659	0.6800
	Backbone+DDL	0.5907	0.0653	0.2310
	Backbone+RWSA	0.6064	0.0641	0.2114
	Backbone+DDL+RWSA	<b>0.5781</b>	<b>0.0624</b>	<b>0.1290</b>
IID-Net	Backbone	0.6596	0.0933	0.5785
	Backbone+DDL	0.6339	0.0791	0.0964
	Backbone+RWSA	0.6420	0.0814	0.1297
	Backbone+DDL+RWSA	<b>0.6181</b>	<b>0.0699</b>	<b>0.0283</b>

#### 4.4 Generality of RWSA

Based on considerations regarding anti-forensics and the requirements stemming from domain adaptation, we meticulously design the RWSA module. It is natural to elegantly insert the RWSA in any existing networks to enhance their anti-forensic performance. To substantiate this assertion, we will quantitatively compare the original versions and updated versions of some state-of-the-art methods, while keeping the experimental setup consistent with Section 4.3. To be precise, we conduct experiments on EdgeConnect [29] and AOT-GAN [38]. The training strategy also adopts their original setting. In line with our approach, we incorporate RWSA into the downsampling layer of the inpainting network. The corresponding numerical results are shown in Table 6. It is evident that RWSA can serve as an effective supplement to address insufficient resistance of traditional inpainting methods against detection.

**Table 6: Anti-forensic performance evaluation of existing methods and these methods with our RWSA module conducted on three detectors.**

Detectors	Model	AUC↓	F1↓	ACC↓
PSCC-Net	EdgeConnect [29]	0.7790	0.1437	0.3890
	EdgeConnect+RWSA	0.6192	0.0666	0.1706
	AOT-GAN [38]	0.8067	0.2487	0.5528
	AOT-GAN+RWSA	0.7590	0.1110	0.2338
Trufor	EdgeConnect	0.8275	0.3627	0.4304
	EdgeConnect+RWSA	0.7066	0.1230	0.1552
	AOT-GAN	0.8232	0.3923	0.5622
	AOT-GAN+RWSA	0.8159	0.3847	0.5478
IID-Net	EdgeConnect	0.7845	0.1582	0.5090
	EdgeConnect+RWSA	0.7621	0.1387	0.4389
	AOT-GAN	0.8617	0.3156	0.7913
	AOT-GAN+RWSA	0.8206	0.2713	0.7162

## 5 Conclusions

In this work, we propose the first end-to-end learning architecture called SafePaint for anti-forensic inpainting, which is achieved with the help of domain adaptation. The task-decoupled strategy ensures that each stage of the network performs its designated role, while domain adaptation achieves compatibility between the background and the foreground of an image. Moreover, the region-wise separated attention module effectively addresses the shortcomings of the previous attention mechanism module and can be naturally integrated as a plug-in into existing inpainting methods to improve their anti-forensic capabilities. Extensive experiments have demonstrated that our SafePaint has very strong anti-forensic ability while maintaining commendable inpainting performance.

## Acknowledgments

This work is supported by National Key R&D Program of China (Grant No. 2022YFB3103500), National Natural Science Foundation of China (Grant Nos. U22A2030, 61972142), Hunan Provincial Funds for Distinguished Young Scholars (Grant No. 2024JJ2025).

## References

- [1] Ivan Anokhin, Pavel Solovev, Denis Korzhenkov, Alexey Kharlamov, Taras Khakhulin, Aleksei Silvestrov, Sergey Nikolenko, Victor Lempitsky, and Gleb Sterkin. 2020. High-resolution daytime translation without domain labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7488–7497.
- [2] Coloma Ballester, Marcelo Bertalmio, Vicente Caselles, Guillermo Sapiro, and Joan Verdera. 2001. Filling-in by joint interpolation of vector fields and gray levels. *IEEE Transactions on Image Processing* 10, 8 (2001), 1200–1211.
- [3] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. 2009. PatchMatch: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics* 28, 3 (2009), 24.
- [4] Marcelo Bertalmio, Guillermo Sapiro, Vincent Caselles, and Coloma Ballester. 2000. Image inpainting. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques*. 417–424.
- [5] Pierre Buysens, Maxime Daisy, David Tschumperlé, and Olivier Lézoray. 2015. Exemplar-based inpainting: Technical review and new heuristics for better geometric reconstructions. *IEEE Transactions on Image Processing* 24, 6 (2015), 1809–1824.
- [6] Antonio Criminisi, Patrick Perez, and Kentaro Toyama. 2003. Object removal by exemplar-based inpainting. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Vol. 2. II–II.
- [7] Soheil Darabi, Eli Shechtman, Connelly Barnes, Dan B Goldman, and Pradeep Sen. 2012. Image melding: Combining inconsistent images using patch-based synthesis. *ACM Transactions on Graphics* 31, 4 (2012), 1–10.
- [8] Ding Ding, Sundaresh Ram, and Jeffrey J Rodriguez. 2018. Image inpainting using nonlocal texture matching and nonlinear filtering. *IEEE Transactions on Image Processing* 28, 4 (2018), 1705–1719.
- [9] Carl Doersch, Saurabh Singh, Abhinav Gupta, Josef Sivic, and Alexei A Efros. 2015. What makes paris look like paris? *Commun. ACM* 58, 12 (2015), 103–110.
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*. 2672–2680.
- [11] Fabrizio Guillaro, Davide Cozzolino, Avneesh Sud, Nicholas Dufour, and Luisa Verdoliva. 2023. Trufor: Leveraging all-round clues for trustworthy image forgery detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 20606–20615.
- [12] Xiefan Guo, Hongyu Yang, and Di Huang. 2021. Image inpainting via conditional texture and structure dual generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 14134–14143.
- [13] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9, 8 (1997), 1735–1780.
- [14] Jie Hu, Li Shen, and Gang Sun. 2018. Squeeze-and-excitation networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7132–7141.
- [15] Jia-Bin Huang, Sing Bing Kang, Narendra Ahuja, and Johannes Kopf. 2014. Image completion using planar structure guidance. *ACM Transactions on Graphics* 33, 4 (2014), 1–10.
- [16] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. 2017. Globally and locally consistent image completion. *ACM Transactions on Graphics* 36, 4 (2017), 1–14.
- [17] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1125–1134.
- [18] Jitesh Jain, Yuqian Zhou, Ning Yu, and Humphrey Shi. 2023. Keys to better image inpainting: Structure and texture go hand in hand. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 208–217.
- [19] Somi Jeong, Youngjung Kim, Eungbeon Lee, and Kwanghoon Sohn. 2021. Memory-guided unsupervised image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6558–6567.
- [20] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016. Perceptual losses for real-time style transfer and super-resolution. In *Proceedings of the European Conference on Computer Vision*. Springer, 694–711.
- [21] Haodong Li and Jiwu Huang. 2019. Localization of deep inpainting using high-pass fully convolutional network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 8301–8310.
- [22] Haodong Li, Weiqi Luo, and Jiwu Huang. 2017. Localization of diffusion-based inpainting in digital images. *IEEE Transactions on Information Forensics and Security* 12, 12 (2017), 3050–3064.
- [23] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. 2018. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European Conference on Computer Vision*. 85–100.
- [24] Hongyu Liu, Bin Jiang, Yi Xiao, and Chao Yang. 2019. Coherent semantic attention for image inpainting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4170–4179.
- [25] Xiaohong Liu, Yaojie Liu, Jun Chen, and Xiaoming Liu. 2022. Psc-net: Progressive spatio-channel correlation network for image manipulation detection and localization. *IEEE Transactions on Circuits and Systems for Video Technology* 32, 11 (2022), 7505–7517.
- [26] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. 2015. Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision*. 3730–3738.
- [27] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. 2022. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11461–11471.
- [28] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. 2018. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957* (2018).
- [29] Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Qureshi, and Mehran Ebrahimi. 2019. Edgeconnect: Structure guided image inpainting using edge prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*. 0–0.
- [30] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. 2016. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2536–2544.
- [31] Weize Quan, Ruisong Zhang, Yong Zhang, Zhifeng Li, Jue Wang, and Dong-Ming Yan. 2022. Image inpainting with local and global refinement. *IEEE Transactions on Image Processing* 31 (2022), 2405–2420.
- [32] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. 2022. Resolution-robust large mask inpainting with fourier convolutions. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 2149–2159.
- [33] D Tschumperle and R Deriche. 2005. Vector-valued image regularization with PDEs: A common framework for different applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 27, 4 (2005), 506–517.
- [34] Haiwei Wu and Jiantao Zhou. 2021. IID-Net: Image inpainting detection network via neural architecture search and attention. *IEEE Transactions on Circuits and Systems for Video Technology* 32, 3 (2021), 1172–1185.
- [35] Qiong Wu, Shao-jie Sun, Wei Zhu, Guo-Hui Li, and Dan Tu. 2008. Detection of digital doctoring in exemplar-based inpainted images. In *International Conference on Machine Learning and Cybernetics*, Vol. 3. IEEE, 1222–1226.
- [36] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. 2018. Generative image inpainting with contextual attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5505–5514.
- [37] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. 2019. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4471–4480.
- [38] Yanhong Zeng, Jianlong Fu, Hongyang Chao, and Baining Guo. 2023. Aggregated Contextual Transformations for High-Resolution Image Inpainting. *IEEE Transactions on Visualization and Computer Graphics* 29, 7 (2023), 3266–3280.
- [39] Yu Zeng, Zhe Lin, Huchuan Lu, and Vishal M Patel. 2021. Cr-fill: Generative image inpainting with auxiliary contextual reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 14164–14173.
- [40] Yu Zeng, Zhe Lin, Jimei Yang, Jianming Zhang, Eli Shechtman, and Huchuan Lu. 2020. High-resolution image inpainting with iterative confidence feedback and guided upsampling. In *Proceedings of the European Conference on Computer Vision*. Springer, 1–17.
- [41] Haoran Zhang, Zhenzhen Hu, Changzhi Luo, Wangmeng Zuo, and Meng Wang. 2018. Semantic image inpainting with progressive generative networks. In *Proceedings of the ACM International Conference on Multimedia*. 1939–1947.
- [42] Lingzhi Zhang, Yuqian Zhou, Connelly Barnes, Sohrab Amirghodsi, Zhe Lin, Eli Shechtman, and Jianbo Shi. 2022. Perceptual artifacts localization for inpainting. In *Proceedings of the European Conference on Computer Vision*. Springer, 146–164.
- [43] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. 2017. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 40, 6 (2017), 1452–1464.
- [44] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*. 2223–2232.
- [45] Manyu Zhu, Dongliang He, Xin Li, Chao Li, Fu Li, Xiao Liu, Errui Ding, and Zhaoxiang Zhang. 2021. Image inpainting by end-to-end cascaded refinement with mask awareness. *IEEE Transactions on Image Processing* 30 (2021), 4855–4866.