

InpDiffusion: Image Inpainting Localization via Conditional Diffusion Models

Kai Wang, Shaozhang Niu*, Qixian Hao, Jiwei Zhang

Beijing Key Lab of Intelligent Telecommunication Software and Multimedia, Beijing University of Posts and Telecommunications, Beijing 100876, China
{kaiwang, szniu, qixianhao123, jwzhang666}@bupt.edu.cn

Abstract

As artificial intelligence advances rapidly, particularly with the advent of GANs and diffusion models, the accuracy of Image Inpainting Localization (IIL) has become increasingly challenging. Current IIL methods face two main challenges: a tendency towards overconfidence, leading to incorrect predictions; and difficulty in detecting subtle tampering boundaries in inpainted images. In response, we propose a new paradigm that treats IIL as a conditional mask generation task utilizing diffusion models. Our method, InpDiffusion, utilizes the denoising process enhanced by the integration of image semantic conditions to progressively refine predictions. During denoising, we employ edge conditions and introduce a novel edge supervision strategy to enhance the model’s perception of edge details in inpainted objects. Balancing the diffusion model’s stochastic sampling with edge supervision of tampered image regions mitigates the risk of incorrect predictions from overconfidence and prevents the loss of subtle boundaries that can result from overly stochastic processes. Furthermore, we propose an innovative Dual-stream Multi-scale Feature Extractor (DMFE) for extracting multi-scale features, enhancing feature representation by considering both semantic and edge conditions of the inpainted images. Extensive experiments across challenging datasets demonstrate that the InpDiffusion significantly outperforms existing state-of-the-art methods in IIL tasks, while also showcasing excellent generalization capabilities and robustness.

Introduction

Image inpainting, essential for digital restoration and other applications, has reached unprecedented realism through advancements in deep learning techniques (Rombach et al. 2022; Zeng et al. 2022; Feng et al. 2023; Park et al. 2024). This sophistication has outpaced traditional tampering detection methods, which depend on visible signs. The field of image forensics now faces the new challenge of detecting subtle, concealed manipulations through the emerging task of Image Inpainting Localization (IIL).

Researchers have responded to this challenge by proposing various strategies, focusing primarily on three design paradigms: 1) Multi-stream frameworks (Dong et al. 2022; Niloy et al. 2023) that leverage multiple input streams to

*Corresponding author.




Figure 1: Visualization results of tampered edges and region predictions captured by InpDiffusion at different sampling stages. \hat{x}_0 and \hat{e} represent the predicted inpainted objects and their edges at different sampling stages, respectively. $\hat{x}_0^{w/o ES}$ are the predictions without edge supervision strategy.

learn comprehensive representations; 2) Multi-stage (Liu et al. 2022; Guo et al. 2023) feature fusion frameworks that integrate features extracted at different levels of the backbone network; 3) Branch frameworks (Xu, Shen, and Lyu 2023; Shi et al. 2023) that implement a single-input, multiple-output architecture, including tasks like word segmentation and auxiliary branches. These techniques are built upon a semantic segmentation foundation, employing learning-based backbones for feature extraction and decoder heads to generate segmentation masks. However, without a sophisticated design, this approach is prone to overconfidence in predictions, leading to incorrect assessments, especially when dealing with increasingly sophisticated inpainting images (Barglazan, Brad, and Constantinescu 2024). The high invisibility of inpainting edges further complicates the task, as it blurs the distinction between the inpainted areas and the original image background, increasing the risk of misidentification.

To address the unique challenges of IIL, we develop a new framework called InpDiffusion. This framework leverages the diffusion model paradigm, known for its strong generative capabilities and sensitivity to conditions. Unlike traditional IIL models, InpDiffusion’s iterative noise reduction mechanism simplifies the refinement process. Moreover, the diffusion model’s stochastic sampling process introduces

variability in predictions, mitigating the risk of overconfidence in the model’s output. However, the direct application of diffusion models to IIL reveals certain limitations, such as limited discriminating power and excessive stochasticity in sampling, along with insufficient mask refinement. To address these limitations, we extract discriminative semantic and edge features from the inpainted images to guide the downstream denoising process. In addition, we introduce a novel edge supervision strategy within the denoising network to enhance InpDiffusion’s ability to perceive subtle edges of locally inpainted objects. InpDiffusion generates both the denoised mask and the denoised edge of the inpainted object at each denoising step. By balancing the loss weights of these dual predictions, we aim to constrain the network and reduce detail loss caused by excessive randomness during each sampling step. More specifically, InpDiffusion integrates an Adaptive Conditional Network (ACN), which simultaneously extracts semantic and edge features from images. The Dual-stream Multi-scale Feature Extractor (DMFE) we designed is the core component of the ACN, enabling the adaptive extraction of multi-scale features to enhance the representation of both semantic and edge features. Then, the extracted semantic feature serves as a guiding clue for the denoising process, while the extracted edge feature is used for edge supervision, thereby strengthening the network’s ability to capture fine details and reducing the randomness of each sample. This innovative dual approach is pivotal for identifying objects with concealed tampering traces, leading to more reliable and consistent results. As shown in Figure 1, when edge supervision is absent, the model treats a small shaded area below the ball as a tampering area due to excessive random sampling. In contrast, the model initially produces rough edges of the tampering region, but with the progress of sampling and continuous constraints of edge features, the model finally gets a refined prediction. InpDiffusion offers several advantages over previous IIL methods: 1) Multiple random sampling of the diffusion model can avoid overconfident point estimation issues. 2) Efficient extraction of multi-scale features allows for a more comprehensive understanding of the intrinsic differences between inpainted objects and the image background. 3) The balance between image semantic supervision and image edge supervision significantly improves the handling of detail loss in inpainted areas.

Our main contributions can be summarized as follows:

- We are the first to treat the IIL task as a mask generation paradigm and use a conditional diffusion framework to form predictions.
- We propose a novel and effective framework called InpDiffusion, which introduces an edge supervision strategy within the denoising network to constrain excessive stochasticity in sampling at each step. Additionally, we develop a Dual-stream Multi-scale Feature Extractor (DMFE) to extract key discriminative features, guiding the downstream denoising processes to produce more accurate and generalized results for IIL tasks.
- Comprehensive experiments confirm that InpDiffusion surpasses current benchmarks in IIL, while exhibiting re-

markable adaptability across various scenarios.

Related Work

Image Inpainting Localization

Recently, models based on CNNs and Transformers have shown impressive results across various visual tasks, leading to an increase in the development of IIL. For example, IID-Net (Wu and Zhou 2021) uses a high-pass filter to extract features from the high-frequency domain. CAT-Net (Kwon et al. 2021) performs double compression detection on JPEG images to obtain an encoder with microscopic feature weights and a parallel combination of macroscopic feature weights to form a dual-stream network. MVSS-Net (Dong et al. 2022) consists primarily of an edge supervision branch and a noise-sensitive branch, using the noise view, the boundary image, and the real image for feature learning. PSCC-Net (Liu et al. 2022) employs a progressive mechanism to predict the operation masks on all scales. It uses a spatio-channel correlation module to guide feature extraction via space attention and channel attenuation. UE-Net (Ji et al. 2023) is trained under dynamic supervision and produces estimated uncertainty maps to refine the detection results of manipulation, which significantly alleviates learning difficulties. TANet (Shi et al. 2023) locates areas for image manipulation using an operator-induced approach, particularly when detecting and segmenting tampered areas that are cleverly embedded in normal images. EditGuard (Zhang et al. 2024) embeds invisible watermarks to simultaneously achieve copyright protection and precise tamper localization, especially for AI-generated content editing methods. Despite progress, many IIL models remain constrained by overconfident point estimates and false segmentation. Moreover, their intricate structures require meticulous calibration to effectively address sophisticated tampering.

Diffusion Models

Diffusion Models, with its two-stage diffusion and reverse process, is a versatile framework for generating and manipulating data, including image generation (Dhariwal and Nichol 2021; Nichol and Dhariwal 2021), inpainting (Chung, Sim, and Ye 2022; Rombach et al. 2022), and editing (Avrahami, Lischinski, and Fried 2022; Choi et al. 2021). Its ability to learn from noise patterns has also made it a strong candidate for tasks such as super-resolution (Li et al. 2022; Wang et al. 2021), deblurring (Whang et al. 2022; Lee et al. 2022), image segmentation (Baranchuk et al. 2021; Bremping et al. 2022) and anomaly detection (Wolleb et al. 2022; Wyatt et al. 2022), where noise analysis is key. Based on diffusion models, our paper introduces the innovative InpDiffusion model, which advances the field of IIL and significantly enhancing the capacity for detailed analysis within this domain.

Method

Motivation. The potential of generative models (Choi et al. 2021; Ho, Jain, and Abbeel 2020) for complex semantic understanding has been extensively explored, offering advanced pattern recognition and iterative refinement ca-

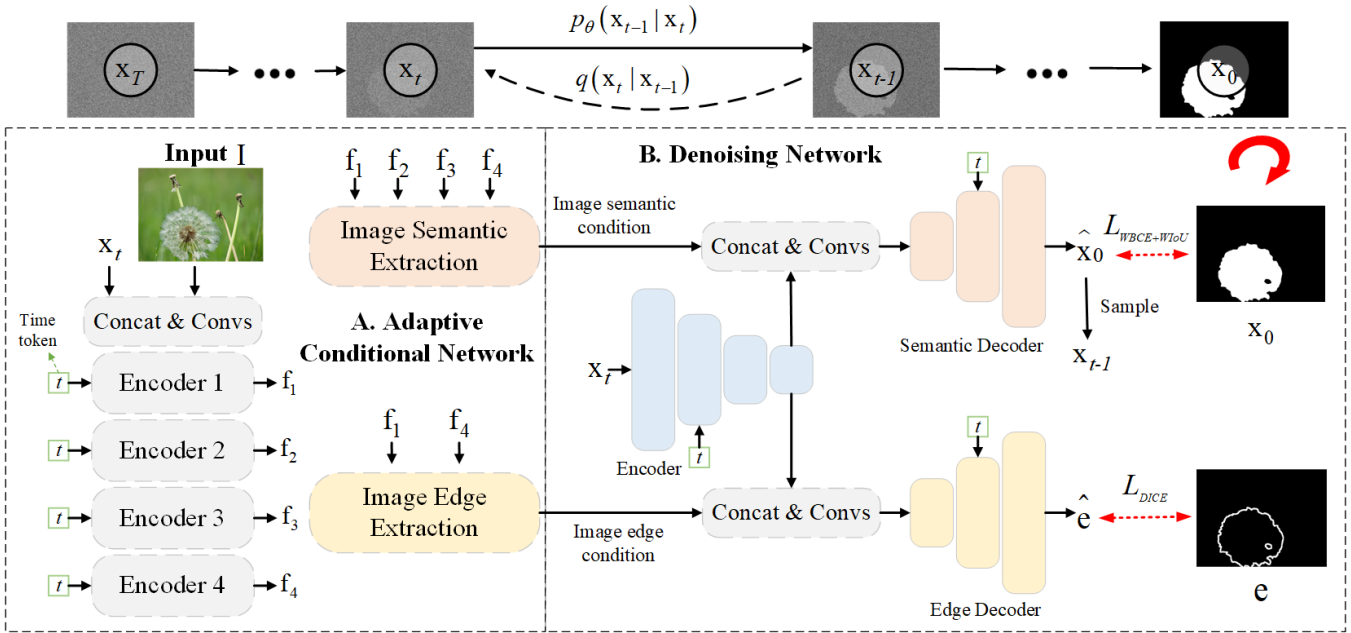


Figure 2: The framework of our InpDiffusion which includes an Adaptive Conditional Network (ACN), and a Denoising Network (DN). Instead of relying on the discriminative learning paradigm, our framework adopts a generative approach to guarantee reliability and generalizability.

pabilities. These capabilities provide a fresh perspective for tackling the increasingly challenging task of IIL. Drawing inspiration from these insights, we are motivated to investigate the potential of diffusion models to further contribute to the IIL field.

Overview. In this section, we introduce our InpDiffusion framework, which progressively generates predictions by using image semantic priors to condition each subsequent step while employing edge priors to constrain the sampling process at each step. As illustrated in Figure 2, our framework comprises two main components: an Adaptive Conditional Network (ACN) and a Denoising Network (DN). We begin by introducing the preliminary, followed by a detailed discussion of the architectures of ACN and DN. Finally, we introduce the loss function used in InpDiffusion.

Preliminary

Conventional IIL relies on a discriminative paradigm, classifying each pixel in an image $I \in \mathbb{R}^{H \times W \times 3}$ against an inpainted map $M \in \mathbb{R}^{H \times W \times 1}$. This approach uses a mapping function F_θ , trained on labeled pairs $\{I_i, M_i\}_{i=1}^N$, to predict the presence ('1') or absence ('0') of inpainted areas. In contrast, our InpDiffusion framework is based on diffusion models, which consist of a forward diffusion process and a reverse denoising process. The forward diffusion process generates noisy masks from clean inputs by iteratively adding noise, following a Markov process. This process can be mathematically described as follows:

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I}), \quad (1)$$

where β_t denotes the variance parameters of the Gaussian noise added at each step t . As t progresses from 1 to T , \mathbf{x}_t evolves from the original inpainted map \mathbf{x}_0 through the process:

$$q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}_t) \mathbf{I}), \quad (2)$$

where $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$, $\alpha_t = 1 - \beta_t$. During the training phase, we modulate the signal-to-noise ratio of the diffusion process by employing an SNR-based variance schedule (Hooeboom, Heek, and Salimans 2023). For more detailed Settings and experiments on this aspect, please refer to *Supplementary*.

In reverse denoising, diffusion models execute a sequence of transitions from $\mathbf{x}_T \rightarrow \mathbf{x}_{T-\Delta} \rightarrow \dots \rightarrow \mathbf{x}_0$, progressively refining initial noise \mathbf{x}_T towards a refined inpainted map \mathbf{x}_0 . Following (Song, Meng, and Ermon 2020), we choose to train a network (InpDiffusion) $f_\theta(\mathbf{x}_t, \mathcal{I}, t)$ to directly estimate the denoised mask $\hat{\mathbf{x}}_0$ conditional on image \mathcal{I} . The network learns the reverse distribution:

$$p(\mathbf{x}_{t-1} | \mathbf{x}_t) := \mathcal{N}(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t, t), \Sigma_\theta(\mathbf{x}_t, t)). \quad (3)$$

Where, $\Sigma_\theta(\mathbf{x}_t, t)$ is set to $\sigma_t^2 = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t$, and $\mu_\theta(\mathbf{x}_t, t)$ can be expressed as:

$$\mu_\theta(\mathbf{x}_t, t) = \frac{\sqrt{\alpha_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} \mathbf{x}_t + \frac{\sqrt{\bar{\alpha}_{t-1}} \beta_t}{1 - \bar{\alpha}_t} \hat{\mathbf{x}}_0. \quad (4)$$

where $\hat{\mathbf{x}}_0$ is predicted by InpDiffusion.

Adaptive Conditional Network (ACN)

The primary function of the Adaptive Convolutional Network (ACN) is to adaptively extract semantic and edge features from inpainted images, allowing the downstream denoising network to recognize inpainted targets and their




Figure 3: Illustration of Dual-stream Multi-scale Feature Extractor (DMFE) .

edges with greater accuracy. The entire ACN consists of two main components: **Hierarchical Feature Extraction** and **Image Semantic and Edge Extraction**.

Hierarchical Feature Extraction. Hidden tampering traces in inpainted images pose significant challenges for feature extraction by the model, which leads to a decline in the mask decoder’s performance. To address this, we input the coarsely predicted mask x_t from the previous step along with the inpainted image \mathcal{I} into the ACN, using x_t as a guiding cue to focus the network on specific regions. Additionally, to enhance feature adaptability across denoising steps, we incorporate the time token (denoted as t) into the feature extraction process. This enables the ACN to autonomously adjust the conditional features according to the temporal stage. As shown in Figure 2 A., we use the Pyramid Vision Transformer (PVT) (Wang et al. 2022) as the backbone of the ACN to extract multi-scale features ($\{f_i\}_{i=1}^4$) from concatenated \mathcal{I} and x_t . Thanks to the Transformer’s token-based design, we can treat time as a token and incorporate it into each layer of patch sequences obtained from the backbone, which are then fed into the corresponding Transformer encoders. Formally, $f_i = R^{-1}(\text{FN}(\text{MA}([t; f_i])))$. Here, $[\cdot; \cdot]$ refers to the concatenation operation, and R^{-1} converts tokens to feature maps. MA represents multi-head self-attention. FN denotes the feedforward neural network. $\{f_i\}_{i=1}^4$ are the hierarchical features we end up with.

Image Semantic and Edge Extraction. The extraction of key discriminative features is essential for denoising networks to accurately identify inpainted targets. To enhance this process, we propose the **Dual-stream Multi-scale Feature Extractor (DMFE)** module as a core component for adaptively extracting image semantic and edge conditions, enabling the capture of additional multi-scale features that the backbone network may not extract. The illustration of




Figure 4: Illustration of Image Semantic and Edge Extraction .

Image Semantic and Edge Extraction is shown in Figure 4. $\{f_i\}_{i=1}^4$ are used to extract semantic feature layer by layer based on DMFE. For the edge condition, we merge the detail-oriented low-level feature f_1 with the spatially informative high-level feature f_4 , to accurately represent the edge feature associated with the inpainted object.

• **Dual-stream Multi-scale Feature Extractor:** Given that the PVT-based backbone network may not extract abundant context information, we drew inspiration from the Inception module and Res2Net module (Gao et al. 2019) and subsequently developed a two-stream multi-scale feature extractor (DMFE). The DMFE consists of two streams. The first stream incrementally applies dilated convolution with a progressively increasing dilated convolution rate, thereby gradually widening the contextual understanding. In contrast, the second stream uses dilated convolution processing with a gradually decreasing rate, which first captures a broad context and then slowly narrows down on specific details. More specifically, we first reduce the number of channels of the input feature f using a 1×1 convolution to facilitate subsequent processing. In the first stream, we utilize four branches to capture features at different scales, sequentially processed with a dilated convolution set with a dilation rate of $\{3, 5, 7\}$. Each branch is equipped with an asymmetric convolution that matches the size of the dilated convolution, reducing computational effort. The output of each branch is added to the input of the next branch. Concurrently, we use residual connections to enlarge the receptive field successively. The general form of an operation is defined as follows:

$$Dout_k \begin{cases} Cov_{3 \times 1}(Cov_{1 \times 3}(Cov(f))) & k=1 \\ Cov_s(Cov(f) + Dout_{k-1}) & k=2, 3, 4 \end{cases} \quad (5)$$

where k is the branch number, $Dout_k(\cdot)$ represents the output of the k th branch in the first stream, $+$ refers to element addition operation. $Cov_s(\cdot)$ denotes the stacked convolutional layer mentioned above. $Cov_{3 \times 1}(\cdot)$ refers to 3×1 convolution operation, $Cov_{1 \times 3}(\cdot)$ refers to 1×3 convolution operation. In contrast to the first stream, the second stream, in turn, processes the features using a dilated convolution set with a dilation rate of $\{7, 5, 3\}$:

$$Uout_k \begin{cases} Cov_s(Cov(f)) & k=1 \\ Cov_s(Cov(f) + Uout_{k-1}) & k=2, 3 \\ Cov_{3 \times 1}(Cov_{1 \times 3}(Cov(f) + Uout_{k-1})) & k=4 \end{cases} \quad (6)$$

where $Uout_k(\cdot)$ denotes the output of the k th branch in the second stream, $Cov(\cdot)$ refers to 1×1 convolution operation.

After that, we add the output of the four branches for each stream to get the result. Finally, we concatenate the results of the two streams followed by a $Cov_{3 \times 3}$ to obtain the output feature, which is computed as:

$$f^m = Cov_{3 \times 3}([(Add_{k=1}^4(Dout_k), Add_{k=1}^4(Uout_k))]). \quad (7)$$

where $Add_{k=1}^4(\cdot)$ refers to the element addition of all four branches.

Denosing Network (DN)

Traditionally, Denosing Networks (DN) aim to decode the denoised mask predictions \hat{x}_0 and x_{t-1} based on the diffusion paradigm. The random sampling in the denoising network may result in insufficient model sensitivity to the edges of concealed inpainted objects. In contrast, our denoising network simultaneously predicts the edge \hat{e} of inpainted objects while generating denoised masks. Through edge loss supervision, we can balance the denoising process of the noised mask, thereby reducing detail loss caused by excessive randomness at each sampling step. As shown in Figure 2 B., our DN comprises an encoder and two decoders. The encoder extracts the positional information of inpainted objects from the noisy mask. One decoder, informed by the image semantic conditions extracted by the ACN, generates the final denoised mask predictions \hat{x}_0 and x_{t-1} . The other decoder, utilizing the image edge conditions extracted by the ACN, produces the final denoised edge predictions \hat{e} . Notably, we employ adaptive group normalization (Dhariwal and Nichol 2021) in both the encoder and decoder to incorporate the time step information t into the convolutional layers, making our DN sensitive to time step variations. Due to the iterative nature of the denoising process, our encoder and decoders employ a simple U-shaped structure, which has proven effective in achieving satisfactory results.

Loss Function

To ensure the denoising network generates accurate results, we measure the loss between the conditional denoising mask output \hat{x}_0 and the ground truth x_0 , as well as between the conditional denoising edge output \hat{e} and the ground truth e . These losses guide the optimization of InpDiffusion. In our work, we introduce the Weighted Binary Cross-Entropy (WBCE) and Weighted Intersection over Union (WIoU) loss functions (Wei, Wang, and Huang 2020) for supervising the denoising result \hat{x}_0 of each iteration of the network. For edge supervision, the Dice loss function (Xie et al. 2020) is used to handle the imbalance between positive and negative samples. The overall loss function for InpDiffusion is formulated as follows:

$$L_{total} = \lambda L_{WBCE + WIoU}(\hat{x}_0, x_0) + \mu L_{DICE}(\hat{e}, e). \quad (8)$$

where, $L_{WBCE + WIoU}(\cdot)$ refers to the combination of WIoU and WBCE loss functions, $L_{DICE}(\cdot)$ refers to Dice loss function, λ and μ are adjustment parameters. We achieve optimal denoising by adjusting the ratio of λ to μ .

Baseline Models	Inpaint 32K			
	TM	CNN	GAN	DM
MantraNet(Wu et al. 2019)	67.3	77.1	76.5	66.2
MVSSNet(Dong et al. 2022)	73.6	84.3	79.4	70.9
TANet(Shi et al. 2023)	85.1	87.2	82.7	73.1
MFINet(Ren et al. 2023)	84.7	91.1	87.9	75.3
CFLNet(Niloy et al. 2023)	79.0	88.5	83.4	72.6
ECNet(Hao et al. 2024)	86.4	90.5	88.3	78.1
Ours-w/o ES	91.6	93.3	93.7	82.4
Ours	93.7	96.2	95.6	84.3

Table 1: Quantitative results of our method and other baseline methods across different inpainting types in the Inpaint32K dataset. TM, CNN, GAN, and DM represent four different inpainting types: Traditional Methods-Based, Convolutional Neural Network-Based, Generative Adversarial Network-Based, and Diffusion Model-Based, respectively. ES stands for the edge supervision strategy proposed in this paper. Results are in % AUC.

Experiments

In this section, we describe the experiments conducted on five different datasets to evaluate the effectiveness of InpDiffusion. Our experiments are primarily conducted on a large dataset called Inpaint32K (Hao 2024), which contains a total of 32,000 inpainted images. This dataset is divided into four distinct inpainting types, with each type comprising 8,000 images: Traditional Methods-Based, CNN-Based, GAN-Based, and Diffusion Model-Based. Other datasets include an inpainting datasets (DID (Wu and Zhou 2021)), an AIGC dataset (AutoSplice (Jia et al. 2023)), a real-life dataset (IMD (Novozamsky, Mahdian, and Saic 2020)), and a traditional benchmark dataset (Nist (NIST 2016)). We compare our model with six baseline models. MantraNet (Wu et al. 2019), MVSSNet (Dong et al. 2022), TANet (Shi et al. 2023), MFINet (Ren et al. 2023), CFLNet (Niloy et al. 2023), and ECNet (Hao et al. 2024). More details about the datasets and baseline models are provided in *Supplementary*. We conduct separate ablation and comparative experiments for each inpainting type in Inpaint32K dataset. Following this, we pre-train our model using 32K inpainted images from the Inpaint32K dataset and evaluate it on the DID dataset. Notably, the other baseline models are also re-trained using the Inpaint32K, even if some are pre-trained on larger datasets. To further demonstrate the generalization performance of our model, we fine-tune it on other datasets of varying types. It’s important to note that DID is exclusively used as a test set, while the other datasets are divided into training and test sets with a 9:1 ratio. We use the pixel-level Area Under the Curve (AUC) score as the evaluation metric.

Implementation Details

We implement our InpDiffusion based on PyTorch using a single NVIDIA A800 with 80GB memory for both training and inference. For efficient training, the model undergoes a total of 150 training epochs. For optimization, the




Figure 5: Visual comparisons with recent SOTA models in challenging scenarios with different inpainting techniques.

Baseline Models	DID									
	GC	CA	SH	EC	LB	RN	NS	LR	PM	SG
MantraNet (Wu et al. 2019)	53.3	54.2	65.7	70.3	60.9	50.2	61.5	58.6	59.3	62.2
MVSSNet (Dong et al. 2022)	59.5	72.9	81.1	86.0	75.4	68.3	85.5	74.2	79.9	91.8
TANet (Shi et al. 2023)	73.6	90.2	93.1	90.6	94.3	90.8	83.6	84.5	90.7	93.4
MFINet (Ren et al. 2023)	72.5	88.3	90.1	93.3	92.7	88.3	79.2	85.7	87.6	94.1
CFLNet (Niloy et al. 2023)	58.8	84.6	84.3	86.8	86.4	65.6	84.3	78.2	82.5	93.3
ECNet (Hao et al. 2024)	69.5	87.3	91.1	91.8	91.6	86.3	80.1	77.7	85.4	92.3
Ours-w/o ES	77.3	92.1	93.3	94.7	95.2	90.4	82.6	87.5	93.6	94.4
Ours	79.4	94.5	96.1	97.2	98.0	93.2	84.9	91.3	96.7	98.3

Table 2: Quantitative comparison of ten different inpainting techniques on the DID dataset under pre-training settings. Results are presented as % AUC.

AdamW (Loshchilov and Hutter 2019) optimizer was utilized along with a batch size set to 32. To adjust the learning rate, we implemented the cosine strategy with an initial learning rate of 0.001. Notably, we set $T = 10$ for sampling and set SNR Shift to $-2 \log(6)$. In the loss function, we set the ratio of λ to μ as 7:3. The ablation of hyperparameters we put in *Supplementary*.

Experiment Results

Quantitative Evaluations. Table 1 shows the quantitative results of our proposed InpDiffusion compared to six other state-of-the-art (SOTA) methods across four different inpainting types on the Inpaint32K dataset. Impressively, across all types of inpainting, InpDiffusion consistently outperformed other models in AUC score, demonstrating significant superiority even compared to the next best-performing model. Notably, when our proposed edge supervision strategy is integrated into the denoising network, the model’s performance improves significantly across each inpainting type. This further highlights the critical role of edge guidance in accurately locating inpainted images with hidden tampering traces. Table 2 outlines the detection performance of InpDiffusion and other SOTA models on the DID dataset across ten different inpainting techniques, following

pre-training on the Inpaint32K dataset. InpDiffusion consistently achieved superior accuracy in nearly all local inpainting techniques. However, MVSSNet exhibited slightly better performance in detecting images inpainted with the NS technique. We attribute this advantage to MVSSNet’s extra pre-training on a larger dataset of 60K images, which enabled it to capture additional key information relevant to IIL.

Qualitative Evaluation. Figure 5 illustrates a comparative visual analysis of InpDiffusion against six established IIL baselines. In this figure, each row illustrates a typical inpainting scenario, with each inpainted region showing complex topological structures and edges. Compared to the first row, which employs traditional methods, and the second row, which utilizes CNN techniques, the inpainted regions in both cases do not fully form a complete semantic structure, presenting a significant challenge for current IIL models. However, the proposed InpDiffusion model excels in mitigating overly confident missegmentations caused by the random sampling inherent in each denoising step. Furthermore, the model effectively delineates clear boundaries due to its iterative denoising paradigm and edge supervision strategy.

Components		Inpaint 32K			
DMFE	ES	TM	CNN	GAN	DM
×	×	88.4	92.6	91.3	80.1
×	✓	91.9	94.0	93.5	82.8
✓	×	91.6	93.3	93.7	82.4
✓	✓	93.7	96.2	95.6	84.3

Table 3: Ablation of model components on Inpaint32K datasets. Results are in % AUC.

Baseline Models	AutoSplice	IMD	Nist
MantraNet(Wu et al. 2019)	88.2	59.5	75.7
MVSSNet(Dong et al. 2022)	93.0	68.2	94.2
TANet(Shi et al. 2023)	93.6	71.7	93.3
MFINet(Ren et al. 2023)	96.1	73.5	95.8
CFLNet(Niloy et al. 2023)	93.2	71.8	94.6
ECNet(Hao et al. 2024)	95.4	75.3	94.2
Ours	97.3	85.5	96.8

Table 4: Quantitative comparison of the other three different kinds of datasets in the fine-tuning setting. Results are in % AUC.

Ablation Studies. We conduct an ablation study on the individual components of InpDiffusion, with detailed results provided in Table 3. The absence of the DMFE module led to a decline in model performance, resulting in a 2.9% drop in AUC score, particularly on inpainted images based on CNN. This underscores the crucial role of DMFE in extracting discriminative image semantic and edge features. The inclusion of DMFE enhances the feature representation of locally inpainted regions, enabling the downstream denoising network to accurately identify and locate these areas. The lack of edge supervision also caused a decrease in InpDiffusion’s detection performance across all types of locally inpainted images, highlighting the importance of edge supervision in complementing the denoising network’s ability to generate precise prediction maps.

Generalization Performance Evaluation. To further evaluate the generalization performance of InpDiffusion and verify its scalability, we fine-tuned InpDiffusion and other SOTA methods using three different datasets involving other types of tampering. The experimental results, as shown in Table 4, indicate that our model continues to outperform other methods in localization performance. In particular, on the challenging real-life dataset IMD, our approach achieved a 10.2% higher AUC score than the second-best method, representing a significant improvement.

Robustness Evaluation. To verify the robustness of InpDiffusion, we conducted experimental evaluations on two types of image attack methods. The first type includes conventional image processing operations such as Gaussian noise and Gaussian filtering, which result in blurred image details and reduced contrast, creating fuzzy boundaries around the tampered areas and making it difficult for the

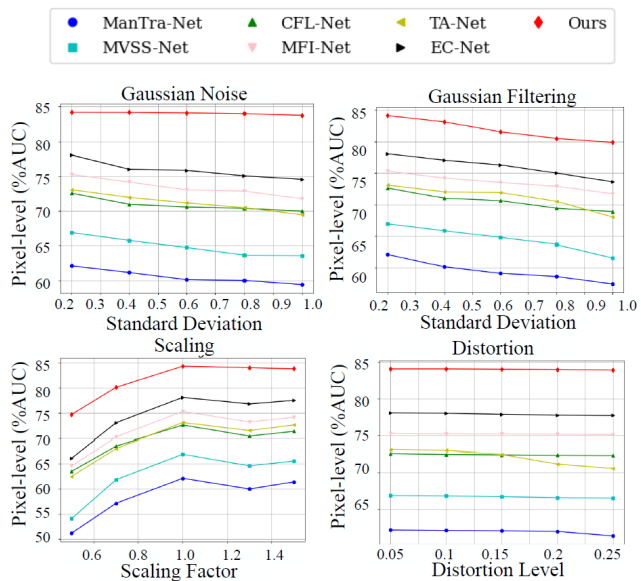


Figure 6: Robustness of seven methods against Gaussian Noise, Gaussian filtering, Scaling and Distortion. The detection is performed on the DM-based inpainted images in Inpaint32K. Our method achieves a substantial lead in IIL performance.

model to accurately locate the tampered regions. The second type involves geometric transformation attacks, such as scaling and distortion, which alter the size, shape, and position of objects in the image, challenging the algorithm’s ability to accurately identify the boundaries of the tampered areas. We tested our model on the relatively difficult-to-detect DM-based inpainted images in Inpaint32K. The experimental results are shown in Figure 6. Our method maintained high robustness against all types of image attacks, particularly against Gaussian noise and image distortion, where InpDiffusion was nearly unaffected by these attacks. We attribute this exceptional performance to InpDiffusion’s use of joint supervision from the inpainted object’s mask and edges to guide the step-by-step denoising process.

Conclusion

In this study, we propose a diffusion-based IIL model named InpDiffusion. InpDiffusion extracts semantic and edge features from inpainted images to guide the denoising process and progressively refine predictions. The introduction of edge supervision helps mitigate the risk of overconfident mispredictions and prevents the loss of subtle boundaries due to excessive randomness. We also introduce an innovative dual stream multiscale feature extractor to capture multi-scale features, enhancing feature representation. Extensive experiments demonstrate that InpDiffusion achieves SOTA performance, showcasing excellent generalization ability and robustness.

Acknowledgments

This work is supported by the National Natural Science Foundation of China(No.61370195) and the Joint Funds of the National Natural Science Foundation of China (No.U1536121).

References

- Avrahami, O.; Lischinski, D.; and Fried, O. 2022. Blended diffusion for text-driven editing of natural images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 18208–18218.
- Baranchuk, D.; Rubachev, I.; Voynov, A.; Khrukov, V.; and Babenko, A. 2021. Label-efficient semantic segmentation with diffusion models. *arXiv preprint arXiv:2112.03126*.
- Barglazan, A.-A.; Brad, R.; and Constantinescu, C. 2024. Image inpainting forgery detection: a review. *Journal of Imaging*, 10(2): 42.
- Brempong, E. A.; Kornblith, S.; Chen, T.; Parmar, N.; Minderer, M.; and Norouzi, M. 2022. Denoising pretraining for semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4175–4186.
- Choi, J.; Kim, S.; Jeong, Y.; Gwon, Y.; and Yoon, S. 2021. Ilvr: Conditioning method for denoising diffusion probabilistic models. *arXiv preprint arXiv:2108.02938*.
- Chung, H.; Sim, B.; and Ye, J. C. 2022. Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12413–12422.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- Dong, C.; Chen, X.; Hu, R.; Cao, J.; and Li, X. 2022. Mvss-net: Multi-view multi-scale supervised networks for image manipulation detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3): 3539–3553.
- Feng, X.; Xu, Y.; Lu, G.; and Pei, W. 2023. Hierarchical Contrastive Learning for Pattern-Generalizable Image Corruption Detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12076–12085.
- Gao, S.-H.; Cheng, M.-M.; Zhao, K.; Zhang, X.-Y.; Yang, M.-H.; and Torr, P. 2019. Res2net: A new multi-scale backbone architecture. *IEEE transactions on pattern analysis and machine intelligence*, 43(2): 652–662.
- Guo, X.; Liu, X.; Ren, Z.; Grosz, S.; Masi, I.; and Liu, X. 2023. Hierarchical fine-grained image forgery detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3155–3165.
- Hao, Q. 2024. Inpaint32K Inpaint 2024 Datasets. https://github.com/QixianHao/Inpaint32K_dataset. [Online; accessed 5 August 2024].
- Hao, Q.; Ren, R.; Wang, K.; Niu, S.; Zhang, J.; and Wang, M. 2024. EC-Net: General image tampering localization network based on edge distribution guidance and contrastive learning. *Knowledge-Based Systems*, 293: 111656.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Hoogeboom, E.; Heek, J.; and Salimans, T. 2023. simple diffusion: End-to-end diffusion for high resolution images. *arXiv preprint arXiv:2301.11093*.
- Ji, K.; Chen, F.; Guo, X.; Xu, Y.; Wang, J.; and Chen, J. 2023. Uncertainty-guided learning for improving image manipulation detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 22456–22465.
- Jia, S.; Huang, M.; Zhou, Z.; Ju, Y.; Cai, J.; and Lyu, S. 2023. Autosplice: A text-prompt manipulated image dataset for media forensics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 893–903.
- Kwon, M.-J.; Yu, I.-J.; Nam, S.-H.; and Lee, H.-K. 2021. CAT-Net: Compression artifact tracing network for detection and localization of image splicing. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 375–384.
- Lee, S.; Chung, H.; Kim, J.; and Ye, J. C. 2022. Progressive deblurring of diffusion models for coarse-to-fine image synthesis. *arXiv preprint arXiv:2207.11192*.
- Li, H.; Yang, Y.; Chang, M.; Chen, S.; Feng, H.; Xu, Z.; Li, Q.; and Chen, Y. 2022. Srdiff: Single image super-resolution with diffusion probabilistic models. *Neurocomputing*, 479: 47–59.
- Liu, X.; Liu, Y.; Chen, J.; and Liu, X. 2022. PSCC-Net: Progressive spatio-channel correlation network for image manipulation detection and localization. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(11): 7505–7517.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled weight decay regularization. In *ICLR*.
- Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, 8162–8171. PMLR.
- Niloy, F. F.; Bhaumik, K. K.; Woo, S. S.; and . 2023. CFL-Net: image forgery localization using contrastive learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 4642–4651.
- NIST. 2016. NIST Nimble 2016 Datasets. <https://www.nist.gov/itl/iad/mig/>. [Online; accessed 2 August 2024].
- Novozamsky, A.; Mahdian, B.; and Saic, S. 2020. IMD2020: A large-scale annotated dataset tailored for detecting manipulated images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision Workshops*, 71–80.
- Park, S.; Moon, S.; Park, S.; and Kim, J. 2024. Localization and manipulation of immoral visual cues for safe text-to-image generation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 4675–4684.
- Ren, R.; Hao, Q.; Niu, S.; Xiong, K.; Zhang, J.; and Wang, M. 2023. MFI-Net: Multi-Feature Fusion Identification Networks for Artificial Intelligence Manipulation. *IEEE Trans-*

- actions on *Circuits and Systems for Video Technology*, 34(2): 1266–1280.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Shi, Z.; Chen, H.; Zhang, D.; and . 2023. Transformer-auxiliary neural networks for image manipulation localization by operator inductions. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(9): 4907–4920.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising Diffusion Implicit Models. In *ICLR*.
- Wang, C.; Yeo, K.; Jin, X.; Coda, A.; Klein, L. J.; and Elmegeeni, B. 2021. S3rp: Self-supervised super-resolution and prediction for advection-diffusion process. *arXiv preprint arXiv:2111.04639*.
- Wang, W.; Xie, E.; Li, X.; Fan, D.-P.; Song, K.; Liang, D.; Lu, T.; Luo, P.; and Shao, L. 2022. Pvt v2: Improved baselines with pyramid vision transformer. *Computational Visual Media*, 8(3): 415–424.
- Wei, J.; Wang, S.; and Huang, Q. 2020. F³Net: fusion, feedback and focus for salient object detection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 12321–12328.
- Whang, J.; Delbracio, M.; Talebi, H.; Saharia, C.; Dimakis, A. G.; and Milanfar, P. 2022. Deblurring via stochastic refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16293–16303.
- Wolleb, J.; Bieder, F.; Sandkühler, R.; and Cattin, P. C. 2022. Diffusion models for medical anomaly detection. In *International Conference on Medical image computing and computer-assisted intervention*, 35–45. Springer.
- Wu, H.; and Zhou, J. 2021. IID-Net: Image inpainting detection network via neural architecture search and attention. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(3): 1172–1185.
- Wu, Y.; AbdAlmageed, W.; Natarajan, P.; and . 2019. Mantra-net: Manipulation tracing network for detection and localization of image forgeries with anomalous features. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9543–9552.
- Wyatt, J.; Leach, A.; Schmon, S. M.; and Willcocks, C. G. 2022. Anoddpm: Anomaly detection with denoising diffusion probabilistic models using simplex noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 650–656.
- Xie, E.; Wang, W.; Wang, W.; Ding, M.; Shen, C.; and Luo, P. 2020. Segmenting transparent objects in the wild. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*, 696–711. Springer.
- Xu, D.; Shen, X.; and Lyu, Y. 2023. UP-Net: Uncertainty-Supervised Parallel Network for Image Manipulation Localization. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(11): 6390–6403.
- Zeng, Y.; Fu, J.; Chao, H.; and Guo, B. 2022. Aggregated contextual transformations for high-resolution image inpainting. *IEEE Transactions on Visualization and Computer Graphics*, 29(7): 3266–3280.
- Zhang, X.; Li, R.; Yu, J.; Xu, Y.; Li, W.; and Zhang, J. 2024. Editguard: Versatile image watermarking for tamper localization and copyright protection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11964–11974.