OpenSDI: Spotting Diffusion-Generated Images in the Open World

Yabin Wang Xi'an Jiaotong University

iamwangyabin@stu.xjtu.edu.cn

Zhiwu Huang* University of Southampton

Zhiwu.Huang@soton.ac.uk

Xiaopeng Hong Harbin Institute of Technology

hongxiaopeng@ieee.org

Abstract

This paper identifies OpenSDI, a challenge for spotting diffusion-generated images in open-world settings. In response to this challenge, we define a new benchmark, the OpenSDI dataset (OpenSDID), which stands out from existing datasets due to its diverse use of large vision-language models that simulate open-world diffusion-based manipulations. Another outstanding feature of OpenSDID is its inclusion of both detection and localization tasks for images manipulated globally and locally by diffusion models. To address the OpenSDI challenge, we propose a Synergizing Pretrained Models (SPM) scheme to build up a mixture of foundation models. This approach exploits a collaboration mechanism with multiple pretrained foundation models to enhance generalization in the OpenSDI context, moving beyond traditional training by synergizing multiple pretrained models through prompting and attending strategies. Building on this scheme, we introduce MaskCLIP, an SPM-based model that aligns Contrastive Language-Image Pre-Training (CLIP) with Masked Autoencoder (MAE). Extensive evaluations on OpenSDID show that MaskCLIP significantly outperforms current state-of-the-art methods for the OpenSDI challenge, achieving remarkable relative improvements of 14.23% in IoU (14.11% in F1) and 2.05% in accuracy (2.38% in F1) compared to the second-best model in localization and detection tasks, respectively. Our dataset and code are available at https://github. com/iamwangyabin/OpenSDI.

1. Introduction

The rise of advanced text-to-image (T2I) generation models, such as Stable Diffusion [9, 41, 45] and other diffusion-based frameworks [27, 47, 61], has profoundly transformed digital content creation, enabling users to produce remarkably lifelike images from simple text prompts. This capability, while expanding creative possibilities, has blurred the boundary between real and AI-generated content, mak-



Figure 1. Three open-world settings of the OpenSDI challenge: (1) user diversity by simulating a range of user preferences, (2) model innovation through the use of multiple advanced diffusion models, and (3) a full manipulation scope that enables both global and local generation.

ing it increasingly difficult to ensure content authenticity. Addressing this challenge requires detectors that can handle three essential open-world settings (Figure 1): (1) the wide range of user preferences (*user diversity*), covering diverse styles, subjects, and creative intentions; (2) the rapid evolution of diffusion models (*model innovation*), introducing significant variability in generated images; and (3) the extensive T2I generation (*manipulation scope*), from a broad spectrum of global image synthesis to precise local changes, object insertions, and background manipulations [46, 62, 63]. We identify this challenge as *Open-world Spotting of Diffusion Images (OpenSDI)*.

In response to the OpenSDI challenge, we introduce OpenSDID, a large-scale dataset specifically designed to address the scarcity of suitable public datasets for benchmarking OpenSDI (Table 1). Our approach incorporates the three essential dimensions of OpenSDI to create a comprehensive and realistic dataset. To reflect *user diversity*, we use multiple large vision-language models (VLMs) [1,

^{*}Corresponding author.

30, 33, 53] to generate a broad range of natural language prompts, simulating how users iteratively refine their creative ideas through diverse, open-world textual descriptions. These prompts are designed to mirror real-world user behavior and preferences, capturing the variety seen in different creative intentions and styles. To address model innovation, we incorporate state-of-the-art (SOTA) T2I diffusion models [9, 27, 41, 45], each with distinct architectures and parameter configurations. This variety in diffusion models allows our dataset to encompass the unique visual characteristics and nuances associated with advanced diffusion models, enhancing the dataset's relevance for evaluating detection and localization across a spectrum of T2I diffusion models. Finally, for manipulation scope, we leverage the semantic understanding of VLMs in combination with Segment Anything Model (SAM) [24] and Florence 2 [60] to translate textual instructions into precise region masks. These masks enable a wide range of manipulations, from local region edits to global image synthesis, simulating the types of modifications that make detection challenging in open-world settings. By combining global and localized manipulations, OpenSDID aims at embodying the complexity required for benchmarking the OpenSDI challenge.

To develop effective approaches for addressing OpenSDI, we face two main barriers: (1) unifying detection and localization frameworks, and (2) generalizing to open-world settings. The unification is challenging due to the high heterogeneity of detection and location. Detection relies on image-level semantics, whereas localization requires pixel-level precision. To tackle these issues, we suggest an approach of building up a mixture of pretrained foundation models. While employing multiple expert models holds significant potential, achieving synergy among them remains a complex challenge. To address this, we introduce a novel framework, Synergizing Pretrained Models (SPM). It is designed to optimize the collaboration between pretrained foundation models for enhanced performance and maintaining generalization capability across both detection and localization tasks in open-world settings. In particular, this approach leverages the strengths of multiple pretrained models through prompt tuning and attention mechanisms, creating a mixture of pretrained foundation models that enhances generalization across diverse real-world scenarios. Building on this SPM approach, we introduce MaskCLIP, a mixture of pretrained models that aligns the capabilities of Contrastive Language-Image Pre-Training (CLIP) with a Masked Autoencoder (MAE) model, outperforming SOTA techniques with considerable relative improvements in both detecting and localizing diffusion-generated image content on the suggested OpenSDI dataset.

In summary, our main contributions are as follows:

• We define the OpenSDI challenge, and introduce a large-

Dataset	# Real	# Fake	User	Model	Full
DiffusionDB [57]	-	14M	√	X	Х
GenImage [64]	1.3M	1.4M	X	✓	X
AutoSplice [21]	2.3K	3.6K	X	X	1
CocoGlide [12]	_	512	X	X	X
HiFi-IFDL [13]	_	1M*	1	X	1
GIM [5]	1.1M	1.1M	X	✓	X
TGIF [38]	3.1K	75K	X	✓	×
OpenSDID	300K	450K	1	✓	1

Table 1. Overview of existing diffusion image datasets and the proposed OpenSDID. User: Simulates diverse preferences. Model: Leverages multiple advanced diffusion models published after 2023. Full: Includes both global and local generation. Note: The asterisk (*) indicates HiFi-IFDL is built from multiple prior datasets. A more comprehensive version is in the suppl. material.

scale dataset (OpenSDID) designed to address the lack of suitable public datasets for OpenSDI.

- We introduce an SPM framework and its derived MaskCLIP to effectively address the OpenSDI challenge.
- Through extensive experiments, MaskCLIP demonstrates markedly superior performance compared to SOTA methods for OpenSDI.

2. Related Works

Diffusion Image Benchmark Datasets. Existing datasets fall short in addressing the OpenSDI challenge due to their significant limitations. Many focus on traditional manipulation techniques such as splicing and copy-move operations. Widely used datasets like Columbia [18], NIST MFC [11], and CASIA [8] concentrate on these conventional methods. HiFi-IFDL [13] created a million-scale dataset that combines various generative models and datasets. However, the manipulation types in HiFi-IFDL are mainly centered around traditional methods and facial editing, and do not involve recent T2I diffusion generators.

To the best of our knowledge, only a few datasets incorporate text-guided editing using diffusion models such as DiffusionDB [57], GenImage [64], CocoGlide [12], AutoSplice [21], TGIF [38] and GIM [5]. DiffusionDB uses a single diffusion model, while CocoGlide and AutoSplice are relatively small in scale and restrict their scope to specific models (Glide [39] and DALL-E2 [43]). While TGIF and GIM offer a broader range of diffusion models, they lack diversity in user preferences as their editing instructions come from a single source. Furthermore, these datasets primarily focus on partial edits rather than comprehensive image generation, limiting their scope. As shown in Table 1, none of these datasets fully addresses all three essential aspects of the OpenSDI challenge.

Detection and Localization Approaches. Early detection methods [3, 54] simply use CNN-based architecture, such

as ResNet [15] and Xception [7], as binary classifiers for distinguishing between real and synthetic images. Several works [10, 35, 49–51] exploit distinctive frequency patterns that differentiate synthetic from real images to improve model generalization performance. However, these methods often struggle with generalization across different types of generators, as they tend to overfit to specific artifacts present in their training data. Recent methods [22, 32, 40, 55, 56] employ large pretrained models, particularly CLIP [42], to construct binary classifiers for deepfake detection. While they evince promising detection performance, they are limited to binary predictions and lack the ability to precisely localize tampered regions.

Another family of approaches (e.g., [19, 58, 59]) provides both detection results and accurate localization of manipulated regions. Early methods like [4, 12, 26, 34] design complex feature extraction and fusion modules across spatial, frequency, and noise domains to enhance the precision of forged image localization. Recently, vision transformer based methods [14, 36, 52] clearly improved the accuracy and efficiency of detection and localization. While these approaches primarily rely on ImageNet pretrained models, newer innovations such as [28, 29] are pushing boundaries by leveraging large-scale vision-language models like CLIP for more robust image forgery detection and localization.

However, most existing methods focus on traditional image manipulations, overlooking those generated globally and locally by advanced T2I diffusion models, leaving a substantial gap in addressing OpenSDI. Furthermore, they have limited generalization when faced with new types of AI-generated images, underscoring a clear limitation in adapting to rapidly evolving image generation technologies.

3. Dataset Creation

We introduce OpenSDID, a large-scale dataset specifically curated for the OpenSDI challenge. Our dataset design addresses the three core requirements essential for open-world spotting of AI-generated content: user diversity, model innovation, and manipulation scope.

User diversity. To simulate the variety of user intentions and styles in real-world settings, we employ an automated pipeline driven by multiple state-of-the-art VLMs, including LLaMA3 Vision [1], LLaVA [33], InternVL2 [6] and Qwen2 VL [53]. These VLMs simulate human-like editing behavior, generating diverse text prompts that mirror the wide range of real-world image manipulation intentions. Based on an authentic set of images from Megalith-10M [2], which is a resource of wholesome, unedited, copyright-free licensed images¹, our pipeline utilizes VLMs to produce varied manipulation instructions. These prompts

emulate common user-driven edits, such as adding or removing objects, altering backgrounds, or changing key image elements. This approach ensures that OpenSDID accurately captures the breadth of user-like diversity in T2I generation, producing a dataset with rich and varied samples that reflect real-world creative diversity.

Model Innovation. As diffusion models evolve rapidly, it is crucial to account for the diverse visual styles and architectural innovations of different T2I diffusion models. To achieve this, we leverage a suite of state-of-the-art generative models-including multiple versions of Stable Diffusion (SD1.5 [45], SD2.1 [45], SDXL [41], SD3 [9]) and Flux.1 [27]. These models span various architectures, parameter configurations, and generation characteristics, providing our dataset with a wide range of model-specific image outputs. By using different models to execute the manipulations specified in the VLM-generated prompts, we ensure that OpenSDID covers a comprehensive spectrum of T2I outputs, reflecting the progression and variety within diffusion-based image synthesis. Our dataset includes distinct manipulation styles and visual attributes unique to each model version, supporting more rigorous testing of detection systems across evolving T2I diffusion models.

Manipulation Scope. As illustrated in Figure 1, we leverages T2I diffusion models for both global image synthesis and precise local edits. To simulate this full range of manipulations, our pipeline integrates SAM [24] and Florence 2 [60] to translate textual prompts into detailed region masks, enabling accurate localized edits. For each image, SAM and Florence generate segmentation masks that delineate specific areas for modifications. During the editing phase, we use the AutoPipelineForInpainting², a standard practice in image edition, to ensure that only the masked regions are modified. To further enhance dataset diversity, we produce multiple variations of each manipulation by using different random seeds. To enhance dataset quality and ensure that generated images align precisely with intended instructions, we employ the CLIP model for optimal image selection. We further diversify our dataset by incorporating traditional editing techniques, such as copy-paste and splicing operations [38], applied at random intervals. Additionally, we include fully synthetic images to ensure comprehensive coverage of diffusion image types. This fusion of global and local generation ensures that OpenSDID captures the complexity of real-world T2I forgeries, covering a comprehensive manipulation spectrum.

The dataset statistics are presented in Table 2, with some samples provided in Figure 2. OpenSDID comprises 300,000 images, evenly distributed between real and fake samples, with 200,000 images (100,000 real from Megalith-10M and 100,000 fake generated by SD1.5) used for train-

¹The images are licensed under: United States Government Work, Public Domain Dedication (CC0), No known copyright restrictions (Flickr commons), or Public Domain Mark.

 $^{^2}$ https://huggingface.co/docs/diffusers/using-diffusers/inpaint



Figure 2. OpenSDID examples of authentic and manipulated images with their corresponding ground-truth masks, and entire generated images.

Model	Traini	ng Set	Test	t Set	Total
Model	Real	Fake	Real	Fake	Images
SD1.5	100K	100K	10K	10K	220K
SD2.1	-	-	10K	10 K	20K
SDXL	-	-	10K	10K	20K
SD3	-	-	10K	10K	20K
Flux.1	-	-	10K	10K	20K
Total	100K	100K	50K	50K	300K

Table 2. Dataset Statistics.

ing, and 100,000 images (20,000 each from SD1.5, SD2.1, SDXL, SD3, and Flux.1) reserved for testing, ensuring a balanced representation across multiple generative models. This design follows common practices in AI-generated image detection [12, 54], evaluating generalization to unseen generators. Additional analysis and details about this dataset, including considerations of ethical and bias issues, are provided in the suppl. material.

4. Proposed Approach

SPM Framework. To tackle the OpenSDI challenge, we focus on two primary issues: (1) unifying detection and localization frameworks, and (2) generalizing to open-world environments. Detection and localization are inherently heterogeneous—the former relies on image-level semantics to classify an image, while the latter demands pixel-level accuracy to identify specific manipulated regions. To address them, we propose a scheme of Synergizing Pretrained Models (SPM) to optimize collaboration among pretrained models, each excelling in learning image- or pixel-level semantics while showing strong generalization capabilities.

The suggested SPM scheme achieves this synergy through two key mechanisms: *prompting* and *attending*. The prompting strategy introduces an additional parametrization to the network input, enabling the model to encode new knowledge without significantly altering the pretrained network components. This approach typically involves minimal fine-tuning or even freezing the original

nal foundation model, ensuring that the pretrained knowledge is largely preserved. The attending strategy, on the other hand, facilitates collaboration between multiple pretrained models via a cross-attention mechanism, aligning feature representations from different models to improve performance. We demonstrate the effectiveness of the SPM scheme with MaskCLIP, which is based on a mixture of CLIP for image-level semantic understanding and MAE for fine-grained spatial precision.

Preliminaries. The CLIP [42] model is a dual-encoder architecture consisting of a text encoder E_t and an vision encoder E_v . The pretrained CLIP has demonstrated impressive zero-shot classification performances by computing cosine similarity between the visual features extracted from images and text features generated from class descriptions.

Given an input image $X \in \mathbb{R}^{H \times W \times 3}$ and text description T, the vision encoder splits the input image X into N image patches $\{x_i\}_{i=1}^N$. Each patch x_i is flattened and projected to obtain visual tokens $v_i \in \mathbb{R}^D$. This set of input visual tokens can be defined as $V_{patch} = \{v_i\}_{i=1}^N$. A learnable [CLS] token $v_{cls} \in \mathbb{R}^D$ is prepended to the sequence of visual tokens to form the complete input sequence, to facilitate a global image representation. Thus the input sequence for layer l is formed by concatenating the CLS token and visual tokens as $V^l = [v_{cls}^l; V_{patch}^l] \in \mathbb{R}^{(N+1) \times D}$.

Through the self-attention of the transformer, the v_{cls} token iteratively aggregates information from all visual tokens, resulting in a global representation of the entire image by the final layer. The [CLS] token v_{cls}^L from the final layer L is linearly projected to get the global image representation g. Similarly, the text encoder embeds the text description T through the transformer to produce text embeddings \mathbf{t} . The image and text feature \mathbf{g} and \mathbf{t} are then normalized and used to compute the cosine similarity between image-text pairs.

MAEs [17] are self-supervised vision transformers that learn rich visual representations through a reconstruction objective of masked image regions. Unlike CLIP, the MAE encoder E_m processes the input image X without a [CLS] token, generating only visual patch tokens $V_m = \{v_{m,i}\}_{i=1}^N$ where $v_{m,i} \in \mathbb{R}^{D_m}$. The decoder attempts to reconstruct the original image patches corresponding to the masked tokens, encouraging the model to learn meaningful features from the spatial context available in the visible tokens.

MaskCLIP Overview. As shown in Figure 3, MaskCLIP integrates the CLIP visual encoder, the CLIP text encoder, and the MAE encoder. Following the MAE study [17], we replace the vanilla MAE decoder with a Feature Pyramid Network (FPN) [31] based decoder, as done in [36], for more precise pixel-level predictions. Based on the SPM scheme, we align the CLIP and MAE components through *prompting* and *attending* to address the OpenSDI challenge. Rather than full fine-tuning, we employ a prompt-tuning strategy [20, 23, 55] with the textual prompting block (Fig-

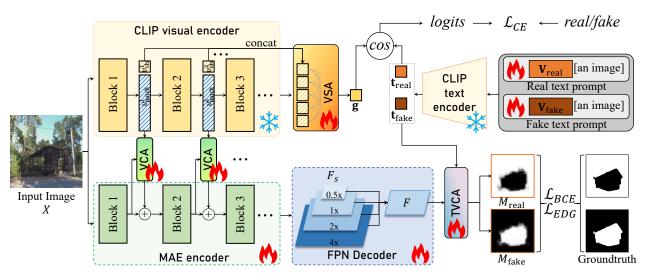


Figure 3. MaskCLIP overview. MaskCLIP inherits core components from pretrained CLIP and MAE: the CLIP vision and text encoders, the MAE encoder, and uses an FPN-style decoder for precise pixel-level predictions. To synergize CLIP and MAE, MaskCLIP introduces one prompting block (top right) and three attention blocks: VCA (visual cross-attention), TVCA (textual-visual cross-attention), and VSA (visual self-attention). Frozen components are marked with snowflakes, while tunable components (marked with flames) are optimized using a balanced objective with \mathcal{L}_{CE} (cross-entropy), \mathcal{L}_{BCE} (binary cross-entropy), and \mathcal{L}_{EDG} (edge-weighted loss).

ure 3), learning a set of continuous prompt vectors for real/fake concepts while preserving CLIP's original parameters to retain its extensive visual-linguistic knowledge. To effectively integrate the prompted CLIP with MAE, the *attending* process is critical. For this purpose, we introduce two types of attending blocks (Figure 3): (1) visual crossattention (VCA), which aligns the CLIP visual encoding features with the MAE encoding features, and (2) textual-visual cross-attention (TVCA), which coordinates the CLIP textual encoding features with the enhanced MAE decoding features. Additionally, we add a visual self-attention (VSA) block to concatenate global features from each CLIP encoding layer for more robust feature extraction. During training, all the prompting and attending blocks (VCA, TVCA, VSA) are jointly optimized.

Prompting. MaskCLIP learns a pair of continuous prompt vectors $\mathbf{V}_c \in \mathbb{R}^{M \times D}$ for $c \in \{\text{real}, \text{fake}\}$, where M denotes the prompt length and D represents the embedding dimension. These learned prompts function analogously to hand-crafted text prompts like "a photo of a <code>[real/fake]</code>", but employ learnable continuous vectors to capture the semantic context more effectively. These learnable prompts \mathbf{V}_c are then processed by the text encoder to generate class-specific embeddings \mathbf{t}_{real} and $\mathbf{t}_{\text{fake}} \in \mathbb{R}^D$, representing the real and fake class concepts, respectively.

Attending. To align CLIP with MAE, MaskCLIP employs the VCA block that combines CLIP's semantic understanding with MAE's spatial reconstruction capabilities. VCA is deployed across multiple network layers to facilitate hierarchical feature fusion between CLIP and MAE. At each

fusion layer l, VCA processes spatial features V_m^l from the MAE encoder alongside the corresponding CLIP visual token features V_{patch}^l . To align these features, we first apply bilinear interpolation to resize V_{patch}^l to the spatial size of MAE features. Subsequently, a 1×1 convolution is employed to project the feature dimension D to D_m . The aligned CLIP features serve as query vectors \mathbf{Q} , while the MAE features V_m^l are projected to form keys and values $\mathbf{K}, \mathbf{V} \in \mathbb{R}^{N\times D_m}$, respectively. We use the multi-head attention mechanism within VCA is computed as:

$$\mathbf{G}^{l} = \operatorname{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^{\top}}{\sqrt{d}}\right)\mathbf{V},\tag{1}$$

where d is the feature dimension for the attention head. To enhance the fusion features, a residual connection is introduced, and the input features for layer l+1 of the MAE is $V_m^{l+1}=V_m^l+\mathbf{G}^l.$

In order to further align the CLIP textual features with the decoding features F, we suggest incorporating a TVCA block to integrate textual information into the decoding process. Given a textual embedding $\mathbf{t}_{\text{real}}, \mathbf{t}_{\text{fake}}$, the decoder performs cross-attention between \mathbf{t} and the decoder's output feature F. First, F is linearly projected to match the dimension of \mathbf{t} . The cross-attention operation uses \mathbf{t} as queries and F as keys and values to get the attention feature map F'. Finally we use a 1×1 convolutional layer to generate the segmentation logits M_{fake} and $M_{\text{real}} \in \mathbb{R}^{H\times W}$.

In addition, the conventional approaches [40, 55] of utilizing only the final layer's [CLS] token v_{cls}^L from the CLIP vision encoder fails to fully leverage the rich hierarchical features learned throughout the network. This limita-

tion is particularly critical for AI-generated image detection, where both low-level artifacts and high-level semantic inconsistencies play crucial roles in discrimination. Thus, rather than relying solely on the last layer output [CLS] token v_{cls}^L of the CLIP vision encoder, we introduce the VSA block to attend [CLS]s from multiple the CLIP encoder layers. VSA leverages a multi-head self-attention mechanism that collects [CLS] token embeddings $v_{cls}^l \in \mathbb{R}^D$ from multiple layers $l \in L$, forming a sequence $V_{cls} = \{v_{cls}^l\}_{l \in L}$. The output features are subsequently aggregated via global average pooling and transformed through a learned linear projection to derive the final image representation \mathbf{g} .

Essential Add-ons. Inspired by prior works [13, 17, 36], we employ an FPN-based decoder [16, 31] that processes the MAE encoder's output features V_m^L at multiple scales $(0.5\times, 2\times, 4\times)$, enabling the capture of both fine and coarse details. At each resolution, a 3×3 convolutional layer is applied to extract information. These multi-scale features F_s are then resized to the original spatial shape (H,W), and are concatenated along the channel dimension to form the feature map F.

To train MaskCLIP for localization, we adopt a combination of binary cross-entropy loss \mathcal{L}_{BCE} and edge-weighted loss \mathcal{L}_{EDG} applied to M_{fake} , following previous works [12, 36]. The total loss function incorporates the detection loss (i.e., cross-entropy loss term, \mathcal{L}_{CE}) to optimize the prompts, VSA, VCA, FPN Decoder, TVCA and the MAE encoder jointly. The weights for all loss terms are set to 1.

During inference, the binary detection is performed by the computing of cosine similarities between the image representation ${\bf g}$ and class-specific text embeddings ${\bf t}_{\rm real}, {\bf t}_{\rm fake}.$ In parallel, the TVCA block generates the forgery localization map $M_{\rm fake}$ to predict potentially manipulated regions.

5. Experiments

Implementation Details. The training of MaskCLIP is conducted with batch size 64 and learning rate 1e-4 for 20 epoch using the Adam optimizer. For data augmentation, we adopt a set of standard settings [12, 34, 36, 37], including Gaussian blurring, JPEG compression, random scaling, horizontal and vertical flips, and random cropping to 512×512 . However, for the CLIP encoder, images are resized to 224×224 pixels to maintain compatibility with its pretrained configuration. In the main experiments, we choose the OpenAI ViT-L/14³ in line with the previous works [25, 40, 48]. For the MAE encoder, we use the ImageNet pretrained weights ViT-B/32 as in previous works [36]. We set the prompt length M to be 10 as previous works [44, 55]. All training is conducted on 4 NVIDIA H100 GPUs.

State-of-the-art Methods. These methods include classical image manipulation localization methods, including CAT-Net [26], MVSS-Net [4], PSCC-Net [34] and Tru-For [12]. For transformer-based methods, we also include ObjectFormer [52], IML-ViT [36] and DeCLIP [48]. In addition to these pixel-level based methods, we also include SOTA image-level based (detection) methods, include CN-NDet [54], UniFD [40], NPR [49], GramNet [35], FreqNet [51], and RINE [25]. All the methods were implemented using the IMDL-BenCo framework [37] and refer to their official codes if possible. We trained these methods on the OpenSDID's training data (Table 2), carefully tuning their hyperparameters to ensure optimal performance that matches their reported benchmarks. A brief description of these methods is provided in the supplementary.

Evaluation Metrics. Following standard practices in [12, 37], we evaluate methods using both pixel-level and imagelevel metrics on the OpenSDID's test data (Table 2). For pixel-level evaluation, we employ the F1-score (F1) and Intersection over Union (IoU). At the image level, we assess performance using F1 and accuracy (Acc).

Localization Results. In Table 3, we show the pixellevel localization performance in terms of F1 and IoU metrics. For the in-domain localization, specifically on the SD1.5 subset, MaskCLIP achieves an IoU of 0.6712 and an F1 score of 0.7563, surpassing all other methods. Notably, our method outperforms IML-ViT, the second best, by 0.92% in IoU and 2.73% in F1 score. In terms of the cross-domain generalization, MaskCLIP shows superior performance across various unseen generators. Our method outperforms previous best methods by significant margins: SD2.1 (1.52% IoU, 0.91% F1), SDXL (16.69% IoU, 16.17% F1), SD3 (23.07% IoU, 21.73% F1), and most notably on Flux.1 (44.69% IoU, 42.34% F1). The average performance across all datasets further emphasizes the superiority of MaskCLIP. It achieves an average IoU of 0.4271 and an average F1 score of 0.4941, surpassing the second best by 14.23% in IoU and 14.11% in F1. The markedly superior in-domain and cross-domain performance suggests that MaskCLIP does not overfit to specific content or manipulation patterns but learns generalizable features crucial for detecting and localizing forgeries across diverse open-world scenarios.

Detection Results. Table 4 displays image-level detection performance (F1-score and Accuracy) across various generative models, separating methods focused solely on image-level detection from those addressing both image and pixellevel tasks. The results illustrate that MaskCLIP achieves the best average performance among all evaluated methods, demonstrating both strong in-domain detection and robust cross-domain generalization capabilities. For the in-domain detection MaskCLIP shows competitive performance with the others. However, in terms of the cross-domain general-

³We also explore CLIP variants and alternative MAE models. To ensure reliability, we conduct experiments with varying hyperparameters and multiple runs. More details are provided in the suppl. material.

	SD	1.5	SD	2.1	SD	XL	SI	D 3	Flu	x.1	ΑV	/G
Method	IoU	F1	IoU	F1	IoU	F1	IoU	F1	IoU	F1	IoU	F1
MVSS-Net [4]	0.5785	0.6533	0.4490	0.5176	0.1467	0.1851	0.2692	0.3271	0.0479	0.0636	0.2983	0.3493
CAT-Net [26]	0.6636	0.7480	0.5458	0.6232	0.2550	0.3074	0.3555	0.4207	0.0497	0.0658	0.3739	0.4330
PSCC-Net [34]	0.5470	0.6422	0.3667	0.4479	0.1973	0.2605	0.2926	0.3728	0.0816	0.1156	0.2970	0.3678
ObjectFormer [52]	0.5119	0.6568	0.4739	0.4144	0.0741	0.0984	0.0941	0.1258	0.0529	0.0731	0.2414	0.2737
TruFor [12]	0.6342	0.7100	0.5467	0.6188	0.2655	0.3185	0.3229	0.3852	0.0760	0.0970	0.3691	0.4259
DeCLIP [48]	0.3718	0.4344	0.3569	0.4187	0.1459	0.1822	0.2734	0.3344	0.1121	0.1429	0.2520	0.3025
IML-ViT [36]	0.6651	0.7362	0.4479	0.5063	0.2149	0.2597	0.2363	0.2835	0.0611	0.0791	0.3251	0.3730
MaskCLIP	0.6712	0.7563	0.5550	0.6289	0.3098	0.3700	0.4375	0.5121	0.1622	0.2034	0.4271	0.4941

Table 3. Pixel-level (localization) performance on OpenSDID. Blue Block: in-domain test, Red Block: cross-domain test, **Bold**: best results, <u>Underline</u>: second best results. More results are in the supplementary material.

	SD	1.5	SD	2.1	SD	XL	SI	D3	Flu	x.1	AV	/G
Method	F1	Acc										
CNNDet [54]	0.8460	0.8504	0.7156	0.7594	0.5970	0.6872	0.5627	0.6708	0.3572	0.5757	0.6157	0.7087
GramNet [35]	0.8051	0.8035	0.7401	0.7666	0.6528	0.7076	0.6435	0.7029	0.5200	0.6337	0.6723	0.7229
FreqNet [51]	0.7588	0.7770	0.6097	0.6837	0.5315	0.6402	0.5350	0.6437	0.3847	0.5708	0.5639	0.6631
NPR [49]	0.7941	0.7928	0.8167	0.8184	0.7212	0.7428	0.7343	0.7547	0.6762	0.7136	0.7485	0.7645
UniFD [40]	0.7745	0.7760	0.8062	0.8192	0.7074	0.7483	0.7109	0.7517	0.6110	0.6906	0.7220	0.7572
RINE [25]	0.9108	0.9098	0.8747	0.8812	0.7343	0.7876	0.7205	0.7678	0.5586	0.6702	0.7598	0.8033
MVSS-Net [4]	0.9347	0.9365	0.7927	0.8233	0.5985	0.7042	0.6280	0.7213	0.2759	0.5678	0.6460	0.7506
CAT-Net [26]	0.9615	0.9615	0.7932	0.8246	0.6476	0.7334	0.6526	0.7361	0.2266	0.5526	0.6563	0.7616
PSCC-Net [34]	0.9607	0.9614	0.7685	0.8094	0.5570	0.6881	0.5978	0.7089	0.5177	0.6704	0.6803	0.7676
ObjectFormer [52]	0.7172	0.7522	0.6679	0.7255	0.4919	0.6292	0.4832	0.6254	0.3792	0.5805	0.5479	0.6626
TruFor [12]	0.9012	0.9773	0.3593	0.5562	0.5804	0.6641	0.5973	0.6751	0.4912	0.6162	0.5859	0.6978
DeCLIP [48]	0.8068	0.7831	0.8402	0.8277	0.7069	0.7055	0.6993	0.6840	0.5177	0.6561	0.7142	0.7313
IML-ViT [36]	0.9447	0.7573	0.6970	0.6119	0.4098	0.4995	0.4469	0.5125	0.1820	0.4362	0.5361	0.5635
MaskCLIP	0.9264	0.9272	0.8871	0.8945	0.7802	0.8122	0.7307	0.7801	0.5649	0.6850	0.7779	0.8198

Table 4. Image-level (detection) performance on OpenSDID. Blue Block: in-domain test, Red Block: cross-domain test, Top Section: image-level detection methods, Bottom Section: pixel-level detection and localization methods. **Bold**: best results, <u>Underline</u>: second best results. More results are in the supplementary material.

ization, MaskCLIP exhibits clearly better performance than all the other methods. On the SD2.1 and SDXL subsets, MaskCLIP even outperforms the SOTA image-level deep-fake detectors, RINE, by **1.24%** and **4.59%**, respectively. For the Flux.1 test data, while image-level detector NPR shows better performance compared to our multi-task detectors, MaskCLIP achieves the best results among all localization methods. For the average result, MaskCLIP significantly outperforms all other methods by **2.38%** on F1 and by **2.05%** on Accuracy.

Qualitative Comparisons. We show some output localization masks of MaskCLIP and other SOTA approaches in Figure 4. As shown across different generators, our method produces masks with well-defined object boundaries and shapes. For instance, the ground truth comparison demonstrates that our approach successfully captures vehicle shapes (SD2.1), clothing textures (SD3), and object shape (SDXL). In contrast, competing methods often exhibit limitations such as over-segmentation of objects or missed detections, highlighting the robustness of our framework. More results are available in the suppl. materials.

Robustness Analysis. We conducted extensive experiments to evaluate detection robustness under different

image degradation scenarios. Following standard practices [12, 37], we employed two common degradation methods: Gaussian blur and JPEG compression, applying them at various intensity levels. The results are presented in Figure 5, and we report the pixel-level F1 scores. For indomain evaluation (SD1.5), our MaskCLIP (shown in red line) demonstrates superior robustness across both degradation types. The proposed MaskCLIP exhibits strong robustness in cross-domain scenarios (SD3), though with a slightly reduced margin compared to in-domain results. This is particularly evident in the JPEG compression test, where our method maintains steady performance even at lower-quality settings. These comprehensive evaluations demonstrate our method's superior robustness in both indomain and cross-domain scenarios, particularly under challenging image quality conditions.

Ablation Study. To thoroughly evaluate the effectiveness of each component in our proposed method, we conduct comprehensive ablation experiments. Starting with a baseline architecture that combines UniFD's [40] classifier with DeCLIP's [48] decoder, we progressively incorporate our key components. Table 5 presents the quantitative results on the OpenSDID SD1.5 test set. The baseline achieves a

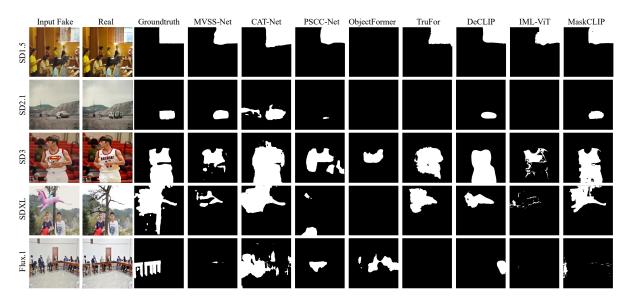


Figure 4. Qualitative results on OpenSDID.

Model Variation	Pixel F1	Image F1
Baseline	0.4290	0.7613
+ Prompt Tuning	0.4485	0.8576
+ VSA	0.4259	0.9083
+ VCA	0.6532	0.9343
+ TVCA (Full MaskCLIP)	0.7563	0.9264

Table 5. Ablation study of the main components in MaskCLIP.

pixel-level F1 of 0.4290 and an image-level F1 of 0.7613. Prompt tuning improves image F1 to 0.8576, showing effective adaptation of CLIP representations. VSA further boosts image-level F1 to 0.9083. The VCA effectively combines CLIP's semantic and MAE's spatial features, which significantly enhances pixel-level F1. Finally, TVCA completes our MaskCLIP model, achieving the highest pixel F1 of 0.7563 while maintaining strong image F1 at 0.9264. These results demonstrate that each component, particularly VCA and TVCA, contributes substantially to the model's overall performance. In the suppl. materials, we replace MAE with SAM [24], achieving a remarkable 0.7873 Pixel F1 on SD1.5, which further demonstrates the scalability of the proposed SPM scheme.

6. Conclusions and Outlook

This paper addresses the OpenSDI challenge of detecting and localizing diffusion-generated images in open-world settings by presenting the OpenSDID benchmark dataset. OpenSDID stands out from existing datasets due to its use of diverse vision-language models and advanced diffusion models to replicate real-world local and global manipulations. To address this challenge, we proposed the approach of Synergizing Pretrained Models (SPM) and introduced

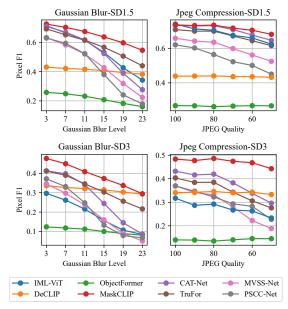


Figure 5. Robustness evaluation of different SOTA methods under image degradation. It compares performance across varying levels of Gaussian Blur (left) and JPEG Compression (right) on both indomain (SD1.5) and cross-domain (SD3) test sets. Results on the rest test data are provided in the suppl. material.

MaskCLIP, an SPM-based model that synergizes CLIP and MAE. This enhances generalization at both image and pixel levels, achieving state-of-the-art detection and localization on OpenSDID. We envision OpenSDID as an open-ended dataset, capable of incorporating images generated by increasingly advanced generative AI models and serving as a standard benchmark for this evolving domain.

7. Acknowledgment

This research was supported by the ECS Fund (504079140) at the University of Southampton, National Natural Science Foundation of China (62376070, 62076195), and the Fundamental Research Funds for the Central Universities (AUGA-5710011522). Yabin Wang was supported by the China Scholarship Council (CSC) Grant No. 202306280130. This work was primarily conducted at the University of Southampton. The authors acknowledge the use of the IRIDIS X High Performance Computing Facility and the Southampton-Wolfson AI Research Machine (SWARM) GPU cluster, funded by the Wolfson Foundation, together with the associated support services at the University of Southampton in the completion of this work.

References

- [1] Meta AI. Llama 3.2-vision: Multimodal large language model. Hugging Face Model Hub, 2024. Model release date: September 25, 2024. 1, 3
- [2] Ollin Boer Bohan. Megalith-10m: A dataset of public domain photographs. https://huggingface.co/datasets/madebyollin/megalith-10m, 2024. Accessed: 2024-10-07. 3
- [3] Lucy Chai, David Bau, Ser-Nam Lim, and Phillip Isola. What makes fake images detectable? understanding properties that generalize. In Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI 16, pages 103–120. Springer, 2020.
- [4] Xinru Chen, Chengbo Dong, Jiaqi Ji, Juan Cao, and Xirong Li. Image manipulation detection by multi-view multi-scale supervision. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 14185–14193, 2021. 3, 6, 7
- [5] Yirui Chen, Xudong Huang, Quan Zhang, Wei Li, Mingjian Zhu, Qiangyu Yan, Simiao Li, Hanting Chen, Hailin Hu, Jie Yang, et al. Gim: A million-scale benchmark for generative image manipulation detection and localization. arXiv preprint arXiv:2406.16531, 2024. 2
- [6] Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhang-wei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, et al. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *Science China Information Sciences*, 67(12):220101, 2024. 3
- [7] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258, 2017. 3
- [8] Jing Dong, Wei Wang, and Tieniu Tan. Casia image tampering detection evaluation database. In 2013 IEEE China summit and international conference on signal and information processing, pages 422–426. IEEE, 2013. 2
- [9] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik

- Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024. 1, 2, 3
- [10] Joel Frank, Thorsten Eisenhofer, Lea Schönherr, Asja Fischer, Dorothea Kolossa, and Thorsten Holz. Leveraging frequency analysis for deep fake image recognition. In *International conference on machine learning*, pages 3247–3258. PMLR, 2020. 3
- [11] Haiying Guan, Mark Kozak, Eric Robertson, Yooyoung Lee, Amy N Yates, Andrew Delgado, Daniel Zhou, Timothee Kheyrkhah, Jeff Smith, and Jonathan Fiscus. Mfc datasets: Large-scale benchmark datasets for media forensic challenge evaluation. In 2019 IEEE Winter Applications of Computer Vision Workshops (WACVW), pages 63–72. IEEE, 2019. 2
- [12] Fabrizio Guillaro, Davide Cozzolino, Avneesh Sud, Nicholas Dufour, and Luisa Verdoliva. 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*, pages 20606–20615, 2023. 2, 3, 4, 6, 7
- [13] Xiao Guo, Xiaohong Liu, Zhiyuan Ren, Steven Grosz, Iacopo Masi, and Xiaoming Liu. Hierarchical fine-grained image forgery detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3155–3165, 2023. 2, 6
- [14] Ruidong Han, Xiaofeng Wang, Ningning Bai, Yihang Wang, Jianpeng Hou, and Jianru Xue. Hdf-net: Capturing homogeny difference features to localize the tampered image. *IEEE Transactions on Pattern Analysis and Machine Intelli*gence, 2024. 3
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 3
- [16] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 6
- [17] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. 4, 6
- [18] Yu-Feng Hsu and Shih-Fu Chang. Detecting image splicing using geometry invariants and camera characteristics consistency. In 2006 IEEE International Conference on Multimedia and Expo, pages 549–552. IEEE, 2006. 2
- [19] Xuefeng Hu, Zhihan Zhang, Zhenye Jiang, Syomantak Chaudhuri, Zhenheng Yang, and Ram Nevatia. Span: Spatial pyramid attention network for image manipulation localization. In Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16, pages 312–328. Springer, 2020. 3
- [20] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *European Conference on Computer Vision*, pages 709–727. Springer, 2022. 4

- [21] Shan Jia, Mingzhen Huang, Zhou Zhou, Yan Ju, Jialing Cai, and Siwei Lyu. Autosplice: A text-prompt manipulated image dataset for media forensics. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 893–903, 2023. 2
- [22] Sohail Ahmed Khan and Duc-Tien Dang-Nguyen. Clipping the deception: Adapting vision-language models for universal deepfake detection. arXiv preprint arXiv:2402.12927, 2024. 3
- [23] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 19113–19122, 2023. 4
- [24] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Con*ference on Computer Vision, pages 4015–4026, 2023. 2, 3, 8
- [25] Christos Koutlis and Symeon Papadopoulos. Leveraging representations from intermediate encoder-blocks for synthetic image detection. In *Computer Vision ECCV 2024*, pages 394–411, Cham, 2025. Springer Nature Switzerland. 6, 7
- [26] Myung-Joon Kwon, Seung-Hun Nam, In-Jae Yu, Heung-Kyu Lee, and Changick Kim. Learning jpeg compression artifacts for image manipulation detection and localization. *In*ternational Journal of Computer Vision, 130(8):1875–1895, 2022. 3, 6, 7
- [27] Black Forest Labs. Flux.1: A 12 billion parameter rectified flow transformer for text-to-image generation. https://github.com/black-forest-labs/flux, 2024.
 A 12 billion parameter rectified flow transformer capable of generating images from text descriptions. 1, 2, 3
- [28] Yingxin Lai, Zhiming Luo, and Zitong Yu. Detect any deep-fakes: Segment anything meets face forgery detection and localization. In *Chinese Conference on Biometric Recognition*, pages 180–190. Springer, 2023. 3
- [29] Dong Li, Jiaying Zhu, Xueyang Fu, Xun Guo, Yidi Liu, Gang Yang, Jiawei Liu, and Zheng-Jun Zha. Noise-assisted prompt learning forimage forgery detection and localization. In *Computer Vision – ECCV 2024*, pages 18–36, Cham, 2025. Springer Nature Switzerland. 3
- [30] Yanghao Li, Haoqi Fan, Ronghang Hu, Christoph Feichtenhofer, and Kaiming He. Scaling language-image pre-training via masking. In *Proceedings of the IEEE/CVF conference* on computer vision and pattern recognition, pages 23390– 23400, 2023. 2
- [31] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 4, 6
- [32] Bo Liu, Fan Yang, Xiuli Bi, Bin Xiao, Weisheng Li, and Xinbo Gao. Detecting generated images by real images. In European Conference on Computer Vision, pages 95–110. Springer, 2022. 3

- [33] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023, 2, 3
- [34] Xiaohong Liu, Yaojie Liu, Jun Chen, and Xiaoming Liu. Pscc-net: Progressive spatio-channel correlation network for image manipulation detection and localization. *IEEE Trans*actions on Circuits and Systems for Video Technology, 32 (11):7505–7517, 2022. 3, 6, 7
- [35] Zhengzhe Liu, Xiaojuan Qi, and Philip HS Torr. Global texture enhancement for fake face detection in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8060–8069, 2020. 3, 6, 7
- [36] Xiaochen Ma, Bo Du, Xianggen Liu, Ahmed Y Al Hammadi, and Jizhe Zhou. Iml-vit: Image manipulation localization by vision transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024. 3, 4, 6, 7
- [37] Xiaochen Ma, Xuekang Zhu, Lei Su, Bo Du, Zhuohang Jiang, Bingkui Tong, Zeyu Lei, Xinyu Yang, Chi-Man Pun, Jiancheng Lv, et al. Imdl-benco: A comprehensive benchmark and codebase for image manipulation detection & localization. arXiv preprint arXiv:2406.10580, 2024. 6, 7
- [38] Hannes Mareen, Dimitrios Karageorgiou, Glenn Van Wallendael, Peter Lambert, and Symeon Papadopoulos. Tgif: Text-guided inpainting forgery dataset. In 2024 IEEE International Workshop on Information Forensics and Security (WIFS), pages 1–6. IEEE, 2024. 2, 3
- [39] Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob Mcgrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In *International Conference on Machine Learning*, pages 16784–16804. PMLR, 2022. 2
- [40] Utkarsh Ojha, Yuheng Li, and Yong Jae Lee. Towards universal fake image detectors that generalize across generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24480–24489, 2023. 3, 5, 6, 7
- [41] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image synthesis. In *The Twelfth International Conference on Learning Representations*, 2024. 1, 2,
- [42] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 3, 4
- [43] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125, 1 (2):3, 2022. 2
- [44] Shuhuai Ren, Aston Zhang, Yi Zhu, Shuai Zhang, Shuai Zheng, Mu Li, Alexander J Smola, and Xu Sun. Prompt pretraining with twenty-thousand classes for open-vocabulary visual recognition. Advances in Neural Information Processing Systems, 36, 2024. 6

- [45] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of* the IEEE/CVF conference on computer vision and pattern recognition, pages 10684–10695, 2022. 1, 2, 3
- [46] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF conference* on computer vision and pattern recognition, pages 22500– 22510, 2023. 1
- [47] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. Advances in neural information processing systems, 35:36479–36494, 2022. 1
- [48] Stefan Smeu, Elisabeta Oneata, and Dan Oneata. Declip: Decoding clip representations for deepfake localization. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2025. 6, 7
- [49] Chuangchuang Tan, Huan Liu, Yao Zhao, Shikui Wei, Guanghua Gu, Ping Liu, and Yunchao Wei. Rethinking the up-sampling operations in cnn-based generative network for generalizable deepfake detection, 2023. 3, 6, 7
- [50] Chuangchuang Tan, Yao Zhao, Shikui Wei, Guanghua Gu, and Yunchao Wei. Learning on gradients: Generalized artifacts representation for gan-generated images detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 12105–12114, 2023.
- [51] Chuangchuang Tan, Yao Zhao, Shikui Wei, Guanghua Gu, Ping Liu, and Yunchao Wei. Frequency-aware deepfake detection: Improving generalizability through frequency space learning, 2024. 3, 6, 7
- [52] Junke Wang, Zuxuan Wu, Jingjing Chen, Xintong Han, Abhinav Shrivastava, Ser-Nam Lim, and Yu-Gang Jiang. Objectformer for image manipulation detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2364–2373, 2022. 3, 6, 7
- [53] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. arXiv preprint arXiv:2409.12191, 2024. 2, 3
- [54] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. Cnn-generated images are surprisingly easy to spot... for now. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8695–8704, 2020. 2, 4, 6, 7
- [55] Yabin Wang, Zhiwu Huang, and Xiaopeng Hong. S-prompts learning with pre-trained transformers: An occam's razor for domain incremental learning. Advances in Neural Information Processing Systems, 35:5682–5695, 2022. 3, 4, 5, 6
- [56] Zhendong Wang, Jianmin Bao, Wengang Zhou, Weilun Wang, Hezhen Hu, Hong Chen, and Houqiang Li. Dire for diffusion-generated image detection. In *Proceedings of the*

- *IEEE/CVF International Conference on Computer Vision*, pages 22445–22455, 2023. 3
- [57] Zijie J Wang, Evan Montoya, David Munechika, Haoyang Yang, Benjamin Hoover, and Duen Horng Chau. Diffusiondb: A large-scale prompt gallery dataset for text-to-image generative models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 893–911, 2023. 2
- [58] Haiwei Wu and Jiantao Zhou. 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, 2021. 3
- [59] Yue Wu, Wael AbdAlmageed, and Premkumar Natarajan. Mantra-net: Manipulation tracing network for detection and localization of image forgeries with anomalous features. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 9543–9552, 2019. 3
- [60] Bin Xiao, Haiping Wu, Weijian Xu, Xiyang Dai, Houdong Hu, Yumao Lu, Michael Zeng, Ce Liu, and Lu Yuan. Florence-2: Advancing a unified representation for a variety of vision tasks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4818– 4829, 2024. 2, 3
- [61] Zeyue Xue, Guanglu Song, Qiushan Guo, Boxiao Liu, Zhuofan Zong, Yu Liu, and Ping Luo. Raphael: Text-to-image generation via large mixture of diffusion paths. Advances in Neural Information Processing Systems, 36, 2024. 1
- [62] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 3836–3847, 2023. 1
- [63] Shihao Zhao, Dongdong Chen, Yen-Chun Chen, Jianmin Bao, Shaozhe Hao, Lu Yuan, and Kwan-Yee K Wong. Uni-controlnet: All-in-one control to text-to-image diffusion models. Advances in Neural Information Processing Systems, 36, 2024. 1
- [64] Mingjian Zhu, Hanting Chen, Qiangyu Yan, Xudong Huang, Guanyu Lin, Wei Li, Zhijun Tu, Hailin Hu, Jie Hu, and Yunhe Wang. Genimage: A million-scale benchmark for detecting ai-generated image. Advances in Neural Information Processing Systems, 36, 2024. 2

OpenSDI: Spotting Diffusion-Generated Images in the Open World

Supplementary Material

In this supplementary material, we provide additional details and results to complement the main paper. Section 1 offers a step-by-step explanation of the OpenSDID creation pipeline, detailing the processes and tools used. Section 2 includes a comprehensive analysis of the dataset's diversity and quality, along with a thorough discussion of the ethical and bias considerations. Section 3 provides a more detailed comparison of the proposed MaskCLIP with related works, highlighting its unique contributions and advantages. To further validate the performance and robustness of MaskCLIP, we present additional experimental results in Section 4 to validate the performance and robustness of MaskCLIP, including cross-dataset evaluations on traditional image forgery datasets and more recent AI-generated image benchmarks, robustness analysis under image degradation, and multiple runs to ensure stability. Additionally, we provide visualizations of the results and datasets in Sections 5 and 6, offering intuitive insights into the data and the method's performance. We are committed to making our research reproducible and accessible to the broader community. Therefore, we have made both the OpenS-DID dataset and the complete codebase publicly available at https://github.com/iamwangyabin/OpenSDI.

1. OpenSDID Creation Pipeline

To create our OpenSDID, we designed a comprehensive framework to automatically generate a large and diverse set of edited images using advanced text-to-image (T2I) diffusion generative models. As illustrated in Figure 1, the process consists of four key steps: (A) load real images from a database, (B) generate textual instructions for editing, (C) create visual masks for editing, and (D) produce AI-generated images with instructions and masks. For global image image generation, OpenSDID merely uses (B) and (D) without using masks to generate images.

We begin by randomly selecting an authentic image from the Megalith-10M database [3]. Next, we randomly choose a state-of-the-art large vision-language model, such as LLaMA3 Vision [1], LLaVA [24], InternVL2 [6], or Qwen2 VL [44], to generate detailed edit suggestions. These models are selected for their ability to produce high-quality and contextually relevant edit instructions, ensuring that the subsequent image manipulations are both creative and realistic. The prompts used to generate these edit suggestions are carefully crafted to guide the model in producing high-quality and relevant modifications. We emphasizes the importance of considering a wide range of potential modifications, including face transformations, hair modifi-

cations, body alterations, clothing and accessory changes, object replacements, background change, style modifications, and other adjustments. Below is an example of the prompts that are likely to be used:

Prompt

You are an image editor with decades of experience in digital manipulation and a reputation for innovation.

Your expertise spans across various genres, including portrait retouching, architectural visualization, product photography, and surreal composites. Your task is to analyze image descriptions and propose compelling, realistic edits that could dramatically enhance or transform the image in unexpected yet believable ways.

When presented with an image, consider a comprehensive range of potential modifications.

- These modifications could be:
 Face transformations
- Hair modifications
- Body alterations
- · Clothing and accessory changes
- · Object replacements or additions
- · Background transformations
- · Architectural style modifications
- · Vehicle transformations
- · Food alterations and adjustments

Your suggestions should be both imaginative and feasible, taking into account the original image's context, composition, and lighting. Strive for a balance between creativity and photorealism, ensuring that your proposed edits could theoretically be executed by a skilled retoucher.

The output text edit instructions specify the selected region in the given real images and the desired content to modify that area. Thus, based on the edit instructions for the previous step, we utilize the Florence-2 [49] and SAM [16]

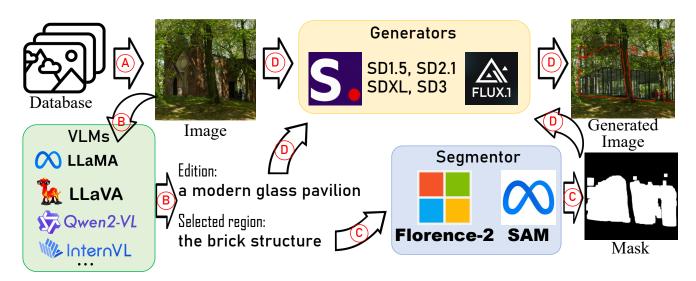


Figure 1. An OpenSDID pipeline for local modification on real image content: (A) Sampling real images from the Megalith-10M dataset, (B) Generating textual instructions for editing using Vision Language Models (VLMs), (C) Creating visual masks for modification through segmentation models, and (D) Producing AI-generated images with image generators based on the instructions and masks. For global image content generation, OpenSDID merely uses (B) and (D) without using real images to produce masks.

models to obtain precise mask regions through open vocabulary detection and segmentation. The Florence-2 model identifies specific regions within the image using text input, while the SAM model refines these regions by generating more precise masks. These masks are further processed by removing small disconnected components and expanded to ensure comprehensive coverage of the areas designated for inpainting.

One of the state-of-the-art T2I diffusion models (SD1.5 [36], SD2.1 [36], SDXL [33], SD3 [10]) and Flux.1 [21]) is then employed to generate new images based on the masked regions and the corresponding prompts. The employed diffusion model takes three inputs: the original image, the generated mask, and the prompt describing the desired new content. To enhance the generation diversity, we randomly adjust key parameters, including the number of inference steps, guidance scale, and strength.

Finally, we use CLIP [34] to compute similarity scores of the generated images and the given edit instructions. Only images that achieve high similarity scores are included in the final dataset, ensuring the maintenance of high-quality standards.

2. OpenSDID Details and Analysis

2.1. Comparison with Existing Datasets

We give a more comprehensive comparison of the existing datasets in Table 1. In terms of user-like diversity, existing datasets have often suffered from limited variability in user input due to their reliance on standardized generation pipelines. This strategy results in reduced diversity across

generated images. While DiffusionDB made progress by collecting generated images from public Discord channels with varied generator parameters, it remained confined to Stable Diffusion 1.5 and globally generated images. Our dataset addresses these limitations by introducing substantial user-like diversity through the integration of multiple state-of-the-art Visual Language Models (VLMs), including LLaMA3 Vision and LLaVA. These VLMs simulate a broad spectrum of human-like editing behaviors, generating diverse text prompts that authentically reflect real-world manipulation intentions. Additionally, we randomize the generation hyperparameters during the diffusion process, including sampling steps, guidance scale, and seed values, to further enhance the diversity of our generated images.

Regarding Model Innovation, many existing datasets, such as DFFD [7] and HiFi-Net [13], primarily focus on images generated by GANs and early diffusion models, lacking representation of recent technological advances. OpenSDID addresses this limitation by incorporating multiple cutting-edge T2I diffusion models, including various versions of Stable Diffusion [10, 33, 36] and Flux.1 [21]. This comprehensive inclusion provides a more robust platform for evaluating detection methods against the rapidly evolving landscape of image synthesis technologies.

With respect to Manipulation Scope, existing datasets often limit themselves to either global or local manipulations. For example, DiffusionDB [46] and GenImage [52] contain only globally synthesized images without local edits, whereas OpenSDID uniquely combines both global and local manipulations. This comprehensive manipulation spectrum more accurately reflects real-world scenarios,

where forgeries frequently involve complex combinations of global and localized alterations.

2.2. Ethical and Bias Issues

The proposed OpenSDID research has been approved by the University's Ethics and Research Governance Online team. Furthermore, we provide a more detailed analysis of potential ethical and bias issues in the dataset.

Copyright Considerations. To ensure the dataset is publicly accessible, we have made every effort to ensure that all locally manipulated images based on [3] are free from copyright restrictions. All images used in this research are freely available under one or more of the following licenses: No Known Copyright Restrictions, United States Government Work, Public Domain Dedication (CC0), or Public Domain Mark. These licenses permit unrestricted use, modification, and distribution of the visual materials. Specifically, the images fall into the following categories: (a) materials free from known copyright restrictions due to their age or provenance, (b) content created by U.S. federal government employees as part of their official duties, (c) works explicitly dedicated to the public domain through Creative Commons Zero designation, or (d) materials marked as public domain due to copyright expiration or other legal provisions. Although these licenses do not require attribution, we have included source citations where applicable to maintain academic integrity and adhere to best practices. The distribution of these copyright statuses is shown in Figure 2.

However, it is important to acknowledge that, given the nature of the Internet, there remains a possibility that a small number of images may have been uploaded by users without proper copyright clearance. While we reserve all rights to the AI-generated images produced in this research, we do not claim ownership of the original source images. Researchers can access these source images independently through the image URLs we have provided in our dataset.

SFW Ensurance. Our dataset consists of images retrieved through the Flickr API, implementing rigorous safety protocols with maximum safety settings (safety_level=1) to ensure content appropriateness [3]. For more comprehensive safety verification, particularly for artificially generated images, we employed multiple state-of-the-art NSFW detectors including:

- GantMAN NSFW Detector [20]
- LAION's CLIP-based NSFW Detector [38]
- Stable Diffusion Safety Checker [35]

Our multi-layered verification process confirmed that 100% of the images maintain SFW status, which can be attributed to both the strict initial filtering and the inherent safety mechanisms in the underlying VLMs and T2I generators.

Potential Biases Analysis. To conduct a comprehensive bias analysis of our dataset, we employed two sophisticated detection frameworks: EasyFace [2] for human de-

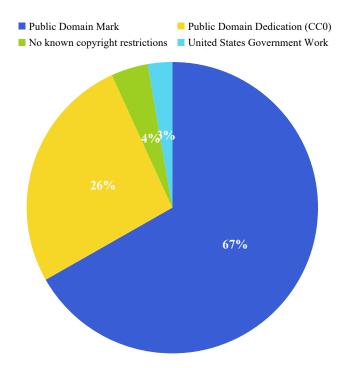


Figure 2. Distribution of copyright statuses for real images in OpenSDID.

mographic analysis and Florence-2 for general object detection.

We utilized the EasyFace detection framework to assess representation across multiple demographic dimensions. Our analysis encompassed binary gender classification, seven distinct racial/ethnic categories, and nine age groups spanning from infancy to elderly, providing a granular view of demographic representation within the dataset. The face detection algorithm was applied across the entire dataset, utilizing pre-trained models optimized for multi-attribute recognition. Table 2 presents the results. While this analysis provides valuable insights into dataset representation, we acknowledge the inherent limitations of automated demographic classification systems, particularly when dealing with edge cases and intersectional identities.

To complement the demographic analysis, we employed Florence-2 for open-world object detection, providing insights into the distribution of general categories within the dataset. The results are visualized in Figure 3.

3. More Related Work Discussion

For training and evaluation, we use the implementation of the IMDL-BenCo framework [28], which includes several state-of-the-art methods for comparison. Here, we briefly describe the methods included in our paper.

CAT-Net (Compression Artifact Tracing Network) [19] is a dual-stream neural network that simultaneously pro-

Dataset	Type	# Real	# Fake	Generator	Adv. T2I	Glo.	Loc.	Users	# Model
UADFV [50]	Face	241	252	GAN	X	X	✓	X	1
DFFD [7]	Face	58K	240K	GAN	X	✓	✓	X	7
FaceForensics++ [37]	Face	1K	4K	GAN	X	X	✓	X	1
DFDC [8]	Face	19K	100K	GAN	X	X	✓	X	2
DeeperForensics [15]	Face	50K	10K	GAN	×	X	1	X	1
CNNSpot [45]	General	362K	362K	GAN	X	✓	X	X	13
GenImage [52]	General	1.3M	1.4M	GAN & Diff.	✓	✓	X	X	8
DiffusionDB [46]	General	_	14M	Diff.	✓	1	X	✓	1
Columbia [30]	General	933	912	Trad.	×	X	1	X	-
CASIA [9]	General	7.2K	5.1K	Trad.	×	X	1	X	-
IMD2020 [31]	General	35K	35K	Trad.	×	X	1	X	_
NIST16 [11]	General	-	564	Trad.	×	X	1	X	_
Coverage [47]	General	100	100	Trad.	×	X	1	X	_
AutoSplice [14]	General	2.3K	3.6K	Diff.	✓	1	1	X	1
CocoGlide [12]	General	_	512	Diff.	✓	X	1	X	1
Dolos [42]	Face	20K	105K	Diff.	✓	1	1	X	4
HiFi-Net [13]*	General	_	1 M	GAN & Diff.	×	1	1	X	10
GIM [5]	General	300K	1.1M	Diff.	✓	X	1	X	3
TGIF [29]	General	3.1K	75K	Diff.	✓	X	✓	X	3
OpenSDID	General	300K	450K	Diff.	✓	/	✓	✓	5

Table 1. Overview of more comprehensive image forgery datasets. "Type" indicates the content category (Face or General). "# Real & # Fake" indicates the number of real and fake images. "Generator" indicates synthesis method type (GAN, Diffusion (Diff.), or Traditional (Trad.)). "Adv. T2I" indicates whether advanced Text-to-Image models (e.g., Stable Diffusion) are used. "Glo." indicates the global image manipulations. "Loc." indicates the local image manipulations. "Users" indicates whether multiple users participate in dataset creation. "# Model" indicates the number of distinct models used for generation. HiFi-Net's locally manipulated images are primarily sourced from previous traditional manipulation datasets and facial deepfake detection datasets.

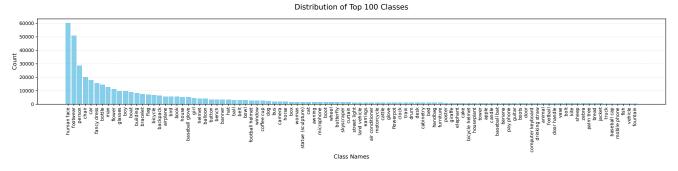


Figure 3. Distribution of top-100 classes in the OpenSDID

cesses both RGB images and JPEG compression artifacts (DCT coefficients) to detect image manipulations. By leveraging both visual content and compression artifacts, it is particularly effective at identifying forgeries, even in compressed images.

MVSS-Net (Multi-View Multi-Scale Supervised Networks) [4] is a dual-branch architecture that simultaneously processes both RGB domain features and noise patterns. The edge-supervised branch utilizes edge residual blocks, while the multi-scale feature learning branch analyzes tampering edge artifacts and noise views of input images.

PSCC-Net (Progressive Spatio-Channel Correlation Network) [25] leverages multi-scale feature learning through dense cross-connections and progressive feature fusion strategies. The model utilizes different sizes of convolutions and perceptual fields to extract valuable information about tampered locations, making it particularly effective at detecting various types of image manipulations.

TruFor [12] uses a Noiseprint++ extractor to process RGB images and obtain learned noise-sensitive fingerprints. These fingerprints, along with the original RGB input, are fed into an encoder that jointly computes features for two

Category	Attribute	Percentage (%)
Gender	Male	53.4
Gender	Female	46.6
	White	46.5
	Black	7.0
	Latino Hispanic	4.4
Race/Ethnicity	East Asian	20.5
	Southeast Asian	2.4
	Indian	1.7
	Middle Eastern	17.5
	0-2	0.4
	3-9	6.4
	10-19	6.7
	20-29	36.4
Age	30-39	18.6
-	40-49	14.2
	50-59	11.6
	60-69	5.0
	70+	0.7

Table 2. Demographic distribution analysis results on the OpenS-DID dataset.

parallel decoding paths: an anomaly decoder for pixel-level forgery localization and a confidence decoder for detection.

ObjectFormer [43] is a transformer-based architecture designed for image manipulation detection and localization. It takes both RGB domain and frequency domain (DCT) as input.

IML-ViT [27] employs a specialized architecture that combines a windowed Vision Transformer (ViT) backbone, which alternates between windowed and global attention blocks to process high-resolution (1024×1024) input images, with a Simple Feature Pyramid Network (SFPN) for multi-scale feature extraction.

DeCLIP [39] leverages CLIP ViT-L/14 as its image encoder and employs a convolutional-based architecture as its decoder. The CLIP image encoder is kept frozen, while only the mask decoder is trained for deepfake localization.

CNNDet [45] is a standard image classifier trained on images generated by a single CNN generator (ProGAN). Through careful data augmentations, it can successfully detect AI-generated images across multiple different architectures and datasets.

UniFD [32] detects AI-generated fake images by extracting features from a frozen CLIP-ViT model and then classifying these features using either nearest neighbor classification or an MLP classifier.

NPR [40] detects synthetic images by analyzing patterns in the relationships between neighboring pixels, which result from upsampling operations in generative networks.

GramNet [26] detects fake faces by analyzing global texture patterns in images, utilizing a specialized neural network that extracts and compares statistical features from real and AI-generated facial textures.

FreqNet [41] integrates a high-frequency representation module and frequency convolutional layer into a lightweight CNN architecture, processing both phase and amplitude spectra between FFT and IFFT operations while forcing continuous focus on high-frequency information to effectively detect deepfakes.

RINE [18] detects AI-generated images by extracting features from multiple intermediate layers of CLIP's Vision Transformer. These extracted features are then used to train a binary classifier that distinguishes between real and fake images.

Methods like CAT-Net and MVSS-Net significantly enhance detection accuracy through their sophisticated integration of multi-domain data, including RGB images, frequency domain information, and noise patterns. Approaches such as ObjectFormer and IML-ViT leverage transformer-based architectures to achieve precise detection and localization capabilities. More recent innovations, including DeCLIP and UniFD, harness the power of pretrained CLIP models for robust feature extraction, pairing them with specialized decoders or classifiers to achieve state-of-the-art detection performance.

The open-world nature of the OpenSDI introduces several challenges. First, the diversity of user preferences and the constant innovation in diffusion models make it difficult to train models that generalize well. Second, the wide range of manipulation scopes, from global image synthesis to local edits, requires models to be robust across different scales and types of modifications.

MaskCLIP advances beyond existing approaches through several key architectural innovations. Unlike methods such as CAT-Net and MVSS-Net that primarily focus on specific artifacts or dual-stream architectures, MaskCLIP leverages a more comprehensive approach through its SPM framework. This framework uniquely combines CLIP's semantic understanding capabilities with MAE's reconstruction power, creating a more robust foundation for detection and localization tasks.

Though some previous studies also leverage pre-trained models, such as CLIP [23, 39] and SAM [22], we integrate both CLIP and MAE to achieve more generalized outcomes. The efficacy of our design is substantiated through comprehensive experimental validation on the OpenSDID dataset, illustrating notable enhancements over current methods in terms of both detection accuracy and localization precision.

4. More Quantitative Results

In this section, in addition to assessing the efficiency and the scalability of MaskCLIP and other state-of-the-art methods, we focus on conducting a series of key additional experiments to examine the performance and capabilities of our proposed method, MaskCLIP, thoroughly across various domains and scenarios. These experiments aim to provide a comprehensive understanding of MaskCLIP's effectiveness in diverse contexts, ranging from traditional image manipulation detection to more advanced challenges posed by AI-generated content.

The main paper focuses on studying the proposed OpenSDI challenge and evaluates crossing the constructed OpenSDI datasets (e.g., SD1.5, Flux.1). In contrast, the evaluation across OpenSDI and non-OpenSDI (other public) datasets is used to study the broader benefits of our OpenSDI dataset and therefore is included in the suppl. material (Tables 5 and 6).

Scalability Analysis. To evaluate the data scalability of MaskCLIP, we conducted experiments by training MaskCLIP and the SOTA method TruFor on different proportions of the OpenSDID dataset. Table 3 presents the pixel-level F1 scores for both methods when trained on 25%, 50%, and 100% of the data. As shown, the performance of MaskCLIP improves as the training data size increases from 25% to 100%, demonstrating its favorable data scaling properties and ability to leverage larger datasets for enhanced performance in OpenSDI tasks.

Data Proportion	TruFor	MaskCLIP
25%	0.5294	0.5737
50%	0.6062	0.5906
100%	0.7100	0.7563

Table 3. Pixel-level F1 scores of MaskCLIP & TruFor (SOTA) on different proportions of the OpenSDID dataset.

Cross-Dataset Evaluation on Traditional Image Forgery Detection and Localization (IMDL) Benchmarks. Table 4 presents a comprehensive comparison between our proposed MaskCLIP method and state-of-the-art image manipulation detection methods across five established forensics datasets (COVERAGE, Columbia, NIST16, CASIAv1, and IMD2020). Following the IMDL-BenCo framework [28], we adopt Protocol-MVSS, where all models are trained exclusively on the CASIAv2 dataset and evaluated directly on other datasets without fine-tuning, enabling a true assessment of zero-shot domain generalization capabilities. The experimental results demonstrate the superior performance of our method across most evaluation scenarios. Notably, our approach achieves the highest average F1 score and outperforms existing methods on three out of five benchmarks. It is important to note that all test benchmarks focus on traditional image manipulation detection tasks, specifically addressing conventional manipulation techniques such as copy-paste and splicing operations. The primary variations across these datasets stem from differences in image content domains and resolutions. But our

OpenSDI challenge tackles a fundamentally different problem: detecting and localizing manipulations generated by advanced AI models, particularly diffusion-based T2I generation methods.

Cross-Dataset Evaluation on Another AI-generated Image Benchmark Dataset (GenImage). To further evaluate the performance of our proposed method, we conduct experiments on the recently introduced GenImage benchmark. This benchmark allows us to compare methods trained on the GenImage SDv1.4 subset with those trained on our OpenSDID. The evaluation is divided into two groups: indataset evaluation and cross-dataset zero-shot evaluation.

In the in-dataset evaluation, the methods are trained and tested on the same distribution. In contrast, the cross-dataset zero-shot evaluation testing models trained on OpenSDID on GenImage dataset. Specifically, GenImage uses ImageNet as its real source data, while OpenSDID is composed of web-collected images, introducing a significant domain gap between the two datasets.

The results in Table 5 demonstrate not only our method's performance but also highlight OpenSDID's superiority in terms of image diversity and realism. The table presents accuracy metrics for various methods across different testing subsets, including Midjourney, SD V1.4, SD V1.5, ADM, GLIDE, Wukong, VQDM, and BigGAN, along with the average accuracy. The results reveal that methods trained on OpenSDID demonstrate superior performance in cross-dataset zero-shot evaluation, despite the domain gap. This suggests that OpenSDID provides a more challenging and diverse training set, enhancing model generalization capabilities. Our proposed method, MaskCLIP, achieves a competitive average accuracy of 77.4%, surpassing several state-of-the-art methods in the cross-dataset setting.

Zero-shot Evaluation on OpenSDID with Different Pretrained Dataset. We conduct comprehensive zero-shot evaluations to assess the generalization capability of various state-of-the-art forgery detection methods across different training settings. Table 6 presents the pixel-level localization performance using IoU and F1 metrics on OpenSDID test sets. We evaluate three representative methods: Tru-For [12], DeCLIP [39], and IML-ViT [27], each tested with their officially released pretrained weights and trained on our OpenSDID.

Several key observations emerge: First, when using original pretrained weights, all methods show limited generalization to T2I diffusion-generated images, with performance declining significantly compared to their reported results on traditional forgery datasets. This indicates that existing methods trained on traditional manipulation data struggle to transfer to diffusion-based forgeries. Second, retraining these methods on OpenSDID substantially improves their performance, particularly for SD1.5 and SD2.1 models. For instance, TruFor's IoU improves from 0.0742

Method	COVERAGE	Columbia	NIST16	CASIAv1	IMD2020	Average
Mantra-Net [48]	0.090	0.243	0.104	0.125	0.055	0.123
MVSS-Net [4]	0.259	0.386	0.246	0.534	0.279	0.341
CAT-Net [19]	0.296	0.584	0.269	0.581	0.273	0.401
ObjectFormer [43]	0.294	0.336	0.173	0.429	0.173	0.281
PSCC-Net [25]	0.231	0.604	0.214	0.378	0.235	0.333
NCL-IML [51]	0.225	0.446	0.260	0.502	0.237	0.334
Trufor [12]	0.419	0.865	0.324	0.721	0.322	0.530
IML-ViT [27]	0.435	0.780	0.331	0.721	0.327	0.519
MaskCLIP	0.451	0.848	0.342	0.725	0.319	0.537

Table 4. Pixel-level F1 (localization) performance across traditional image forensics benchmark datasets. All methods are trained on CASIAv2 and do the cross-dataset evaluation on these benchmarks.

Method				Testing S	Subset				Aug
Meniou	Midjourney	SD V1.4	SD V1.5	ADM	GLIDE	Wukong	VQDM	BigGAN	Acc.(%)
ResNet-50 [†]	54.9	99.9	99.7	53.5	61.9	98.2	56.6	52.0	72.1
DeiT-S [†]	55.6	99.9	99.8	49.8	58.1	98.9	56.9	53.5	71.6
Swin-T [†]	62.1	99.9	99.8	49.8	67.6	99.1	62.3	57.6	74.8
CNNDet [†]	52.8	96.3	95.9	50.1	39.8	78.6	53.4	46.8	64.2
Spec [†]	52.0	99.4	99.2	49.7	49.8	94.8	55.6	49.8	68.8
F3Net [†]	50.1	99.9	99.9	49.9	50.0	99.9	49.9	49.9	68.7
GramNet [†]	54.2	99.2	99.1	50.3	54.6	98.9	50.8	51.7	69.9
PSCC-Net [25]	65.0	94.9	95.3	55.0	56.2	79.5	53.6	59.4	69.9
MVSS-Net [4]	66.7	84.4	85.2	60.3	62.9	69.8	67.8	70.0	70.9
TruFor [12]	55.6	46.0	45.0	66.0	60.8	51.0	70.7	70.4	58.2
DeCLIP [39]	56.3	79.8	79.0	72.4	75.5	79.9	77.5	81.3	75.2
MaskCLIP	51.5	95.0	96.3	70.6	75.5	73.8	78.3	77.8	77.4

Table 5. Image-level (detection accuracy) performance on the GenImage benchmark dataset [52]. The methods are divided into two groups based on their training data: the first group (marked with †) is trained on GenImage SDv1.4, while the second group (no marks) is trained on OpenSDID.

to 0.6342 on SD1.5 after retraining. These results highlight both the challenge and importance of developing robust detection methods specifically designed for diffusion-based forgeries, as traditional datasets show limited effectiveness in this emerging threat landscape.

More Robustness Analysis. Figure 4 shows the performance comparison of different methods under image degradation conditions, specifically Gaussian blur and JPEG compression, across three datasets (SD2.1, SDXL, and Flux.1). Under varying levels of Gaussian blur (3-23) and JPEG compression quality (60-100), MaskCLIP consistently demonstrates superior robustness compared to other approaches. Particularly notable is MaskCLIP's performance on the SD2.1 dataset, where it maintains strong F1 scores even as image quality degrades. While performance naturally decreases with more severe degradation, MaskCLIP exhibits more graceful degradation compared to competing methods, maintaining its lead across different datasets and degradation types. This consistent perfor-

mance advantage highlights the robust nature of our approach in handling various real-world image quality challenges.

Comprehensive Analysis of Pretrained Model Combinations. Since our proposed Synergizing Pretrained Models (SPM) learning scheme leverages multiple pretrained models, with MaskCLIP serving as just one implementation example, we conducted extensive experiments with various model combinations. We explored different CLIP variants (ViT-B/32, OpenCLIP) and alternative pixel-wise encoders to thoroughly evaluate the SPM approach. Table 7 presents a systematic comparison of different encoder combinations. The OpenAI ViT-B/32 paired with MAE-base achieves F1 scores of 0.7227 and 0.4472 on SD1.5 and SD3 datasets, respectively. Scaling up to Open-CLIP ViT-L/14 yields improved performance, demonstrating the advantages of a larger vision transformer architecture. Notably, substituting MAE-base with Dinov2-base results in decreased performance (F1 scores of 0.6278 and

		SD	1.5	SD	2.1	SD	XL	SI	D3	Flu	x.1	AV	/G
Method	Data	IoU	F1										
TruFor [12]	Trufor [†]	0.0742	0.1073	0.0770	0.1115	0.0704	0.1035	0.0996	0.1424	0.1019	0.1464	0.0846	0.1222
TruFor [12]	OpenSDID	0.6342	0.7100	0.5467	0.6188	0.2655	0.3185	0.3229	0.3852	0.0760	0.0970	0.3691	0.4259
DeCLIP [39]	Dolos-LDM	0.0138	0.0218	0.0131	0.0210	0.0089	0.0144	0.0145	0.0230	0.0070	0.0115	0.0115	0.0183
DeCLIP [39]	Dolos-Lama	0.0093	0.0151	0.0098	0.0158	0.0057	0.0098	0.0155	0.0251	0.0048	0.0085	0.0090	0.0149
DeCLIP [39]	Dolos-Pluralistic	0.0145	0.0233	0.0154	0.0245	0.0064	0.0108	0.0182	0.0292	0.0069	0.0116	0.0123	0.0199
DeCLIP [39]	Dolos-Repaint	0.0254	0.0377	0.0221	0.0342	0.0184	0.0284	0.0344	0.0522	0.0162	0.0253	0.0233	0.0356
DeCLIP [39]	OpenSDID	0.3718	0.4344	0.3569	0.4187	0.1459	0.1822	0.2734	0.3344	0.1121	0.1429	0.2520	0.3025
IML-ViT [27]	Trufor [†]	0.0806	0.1143	0.0825	0.1165	0.0746	0.1066	0.1279	0.1750	0.1295	0.1768	0.0990	0.1378
IML-ViT [27]	CASIAv2	0.0248	0.0384	0.0228	0.0366	0.0213	0.0337	0.0266	0.0418	0.0290	0.0460	0.0249	0.0393
IML-ViT [27]	OpenSDID	0.6651	0.7362	0.4479	0.5063	0.2149	0.2597	0.2363	0.2835	0.0611	0.0791	0.3251	0.3730
MaskCLIP	CASIAv2	0.0312	0.0465	0.0289	0.0442	0.0256	0.0398	0.0334	0.0502	0.0358	0.0532	0.0310	0.0468
MaskCLIP	OpenSDID	0.6712	0.7563	0.5550	0.6289	0.3098	0.3700	0.4375	0.5121	0.1622	0.2034	0.4271	0.4941

Table 6. Pixel-level (localization) performance on OpenSDID. Test SOTA methods pretrained weights on various training data and do zero-shot testing on the OpenSDID's test sets. Trufor[†] indicates training on a combined dataset including CASIA v2, FantasticReality [17], IMD2020, and tampered versions of COCO and RAISE datasets [19], which is the training setting of Trufor.

Met	hod		evel F1	
Encoder1	Encoder2	Params.	SD1.5	SD3
OA ViT-B/32	MAE-base	96M	0.7227	0.4472
OC ViT-L/14	MAE-base	114M	0.7363	0.4908
OA ViT-L/14	Dinov2-base	114M	0.6278	0.3521
OA ViT-L/14	MAE-base	114M	0.7563	0.3700
OA ViT-L/14	SAM-base	126M	0.7873	0.5773

Table 7. Performance comparison of different encoder combinations. We evaluate various pretrained models as Encoder1 (CLIP variants) and Encoder2 (pixel-wise encode models). OA and OC denote OpenAI and OpenCLIP respectively. The results show pixel-level F1 scores on both SD1.5 and SD3 datasets. Params. indicates the trainable parameters of the model, where we only keep the CLIP model frozen.

0.3521), suggesting that MAE's self-supervised pretraining approach is more effective for our specific task. The optimal performance is achieved by combining OpenAI ViT-L/14 with SAM-base, reaching F1 scores of 0.7873 and 0.5773. This superior performance indicates that SAM's segment-focused pretraining provides particularly valuable features for our segmentation task, albeit at the cost of increased computational overhead during both training and inference. These comprehensive experiments underscore two key findings: (1) the critical importance of selecting appropriate pretrained weights, and (2) the significant performance benefits that can be achieved through larger model architectures and segment-aware pretraining strategies.

Analysis of Loss Function Components. The loss function in MaskCLIP follows established approaches from previous works [4, 12, 27], incorporating three components: cross-entropy loss (\mathcal{L}_{CE}), binary cross-entropy loss (\mathcal{L}_{BCE}), and edge-weighted loss (\mathcal{L}_{EDG}). While we employ a simple balanced weighting scheme for these loss terms, it is crucial to analyze the impact of each component in our objective function. Table 8 presents experimental results with various

Loss	s Compo	nents	Pixel-level F1		
$\mathcal{L}_{ ext{CE}}$	$\mathcal{L}_{ ext{BCE}}$	$\mathcal{L}_{ ext{EDG}}$	SD1.5	SD3	
1.0	2.0	1.0	0.7575	0.3573	
2.0	1.0	1.0	0.7620	0.3373	
1.0	1.0	0.0	0.7299	0.3108	
1.0	1.0	1.0	0.7563	0.3700	

Table 8. Ablation study on three loss function components. \mathcal{L}_{CE} denotes cross-entropy loss, \mathcal{L}_{BCE} represents binary cross-entropy loss, and \mathcal{L}_{EDG} is the edge-weighted loss.

Method	SD1.5	SD2.1	SDXL	SD3	Flux.1	AVG
TruFor [12]	$0.728_{\pm0.032}$	$0.601_{\pm 0.029}$	$0.289_{\pm 0.059}$	$0.330_{\pm 0.018}$	$0.071_{\pm 0.018}$	$0.404_{\pm0.031}$
MaskCLIP	$0.765_{\pm 0.008}$	$0.628_{\pm 0.021}$	$0.381_{\pm 0.013}$	$0.492_{\pm 0.019}$	$0.188_{\pm 0.009}$	$0.471_{\pm 0.014}$

Table 9. Performance of TruFor and MaskCLIP with multiple runs on OpenSDID.

weight combinations of the three loss terms. Our findings underscore the importance of maintaining a balanced objective function that appropriately weights binary classification and edge preservation.

Multiple Runs Analysis. Although most previous works don't perform multiple runs to study robustness, we conducted three independent training runs for both MaskCLIP and the baseline TruFor method. As shown in Table 9, MaskCLIP exhibits notably smaller standard deviations across all subsets, indicating more stable and reliable performance compared to the baseline.

5. More Qualitative Results

To provide a comprehensive evaluation of our method across different Diffusion models, we present additional qualitative results on various subsets of the OpenSDID dataset. Figures 5, 6, 7, 8, and 9 demonstrate our approach's performance on images generated by SD1.5, SD2.1, SD3, SDXL, and Flux.1 models respectively.

The results showcase our method's effectiveness across various generator models. For the Flux.1 subset (Figure 9),

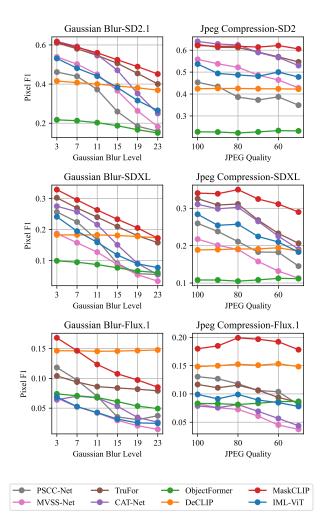


Figure 4. Robustness evaluation of different SOTA methods under image degradation on OpenSDID. It compares performance across varying levels of Gaussian Blur (left) and JPEG Compression (right).

our method exhibits strong zero-shot performance. However, in the interest of transparency, we also present several failure cases (Figure 10) that highlight our method's current limitations. These challenging cases typically involve highly photorealistic images or compositions with unique elements that closely resemble natural photographs. These observations highlight the significant advancement of current AI image generation technology. This rapid development poses higher demands on image discrimination techniques and indicates that future research needs to develop more robust detection methods.

6. Samples of OpenSDID

Figures 11, Figures 12, Figures 13, Figures 14, and Figures 15 present more samples of our datasets.

References

- [1] Meta AI. Llama 3.2-vision: Multimodal large language model. Hugging Face Model Hub, 2024. Model release date: September 25, 2024. 1
- [2] Sithu Aung. Easyface: Easy face analysis tool with sota models. https://github.com/sithu31296/EasyFace, 2022. 3
- [3] Ollin Boer Bohan. Megalith-10m: A dataset of public domain photographs. https://huggingface.co/datasets/madebyollin/megalith-10m, 2024. Accessed: 2024-10-07. 1, 3
- [4] Xinru Chen, Chengbo Dong, Jiaqi Ji, Juan Cao, and Xirong Li. Image manipulation detection by multi-view multi-scale supervision. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 14185–14193, 2021. 4, 7, 8
- [5] Yirui Chen, Xudong Huang, Quan Zhang, Wei Li, Mingjian Zhu, Qiangyu Yan, Simiao Li, Hanting Chen, Hailin Hu, Jie Yang, et al. Gim: A million-scale benchmark for generative image manipulation detection and localization. arXiv preprint arXiv:2406.16531, 2024. 4
- [6] Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhang-wei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, et al. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *Science China Information Sciences*, 67(12):220101, 2024.
- [7] Hao Dang, Feng Liu, Joel Stehouwer, Xiaoming Liu, and Anil K Jain. On the detection of digital face manipulation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern recognition, pages 5781–5790, 2020. 2, 4
- [8] Brian Dolhansky, Joanna Bitton, Ben Pflaum, Jikuo Lu, Russ Howes, Menglin Wang, and Cristian Canton Ferrer. The deepfake detection challenge (dfdc) dataset. arXiv preprint arXiv:2006.07397, 2020. 4
- [9] Jing Dong, Wei Wang, and Tieniu Tan. Casia image tampering detection evaluation database. In 2013 IEEE China summit and international conference on signal and information processing, pages 422–426. IEEE, 2013. 4
- [10] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In Forty-first International Conference on Machine Learning, 2024. 2
- [11] Haiying Guan, Mark Kozak, Eric Robertson, Yooyoung Lee, Amy N Yates, Andrew Delgado, Daniel Zhou, Timothee Kheyrkhah, Jeff Smith, and Jonathan Fiscus. Mfc datasets: Large-scale benchmark datasets for media forensic challenge evaluation. In 2019 IEEE Winter Applications of Computer Vision Workshops (WACVW), pages 63–72. IEEE, 2019. 4
- [12] Fabrizio Guillaro, Davide Cozzolino, Avneesh Sud, Nicholas Dufour, and Luisa Verdoliva. 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*, pages 20606–20615, 2023. 4, 6, 7, 8

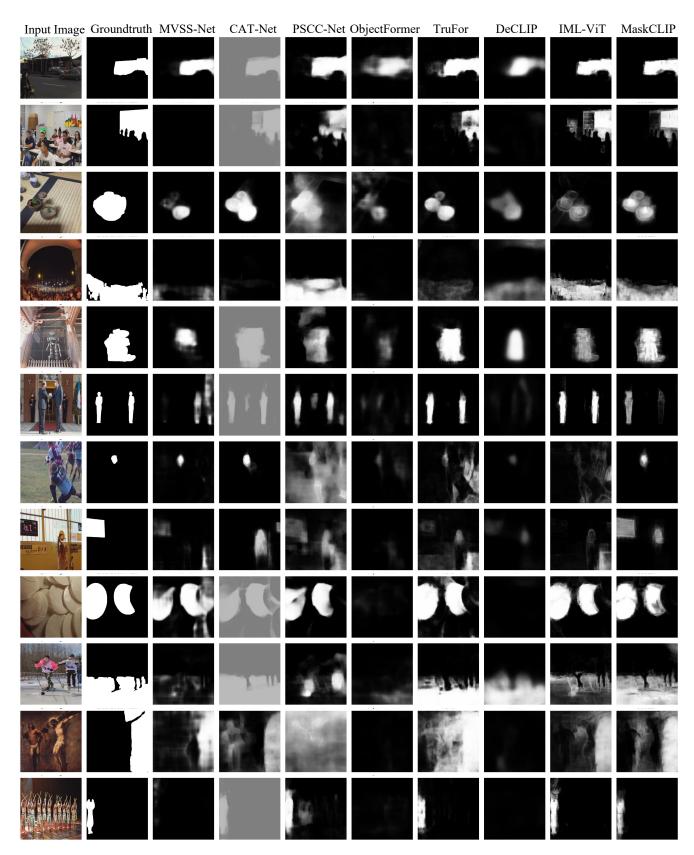


Figure 5. Qualitative results on the OpenSDID SD1.5 testset.



Figure 6. Qualitative results on the OpenSDID SD2.1 test set.



Figure 7. Qualitative results on the OpenSDID SD3 test set.

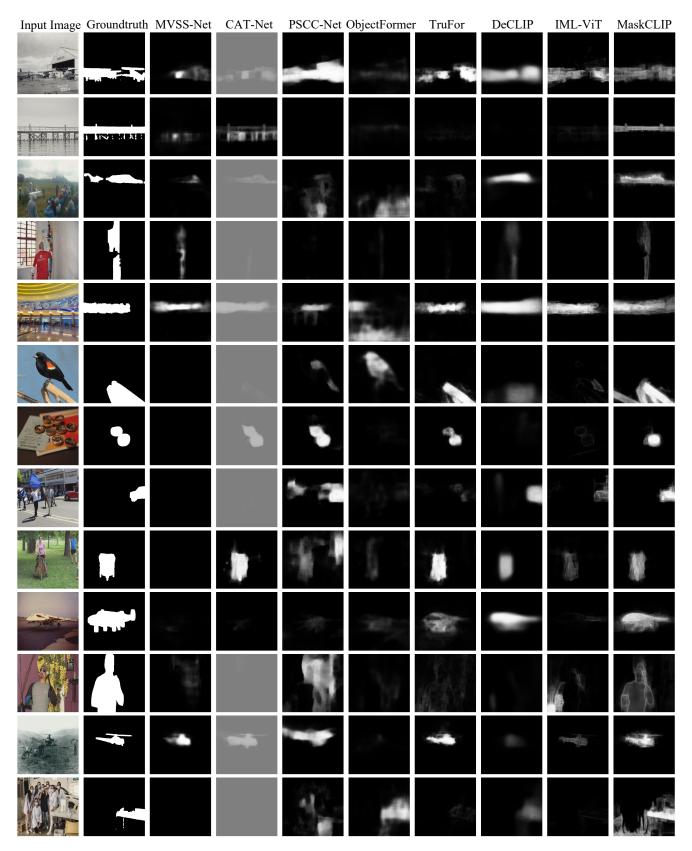


Figure 8. Qualitative results on the OpenSDID SDXL test set.

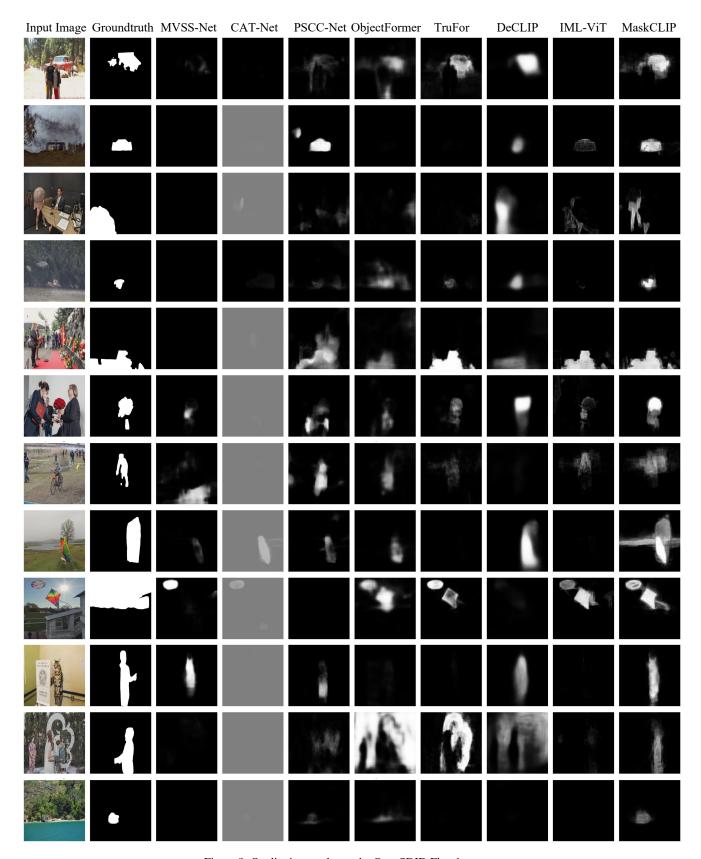


Figure 9. Qualitative results on the OpenSDID Flux.1 test set.

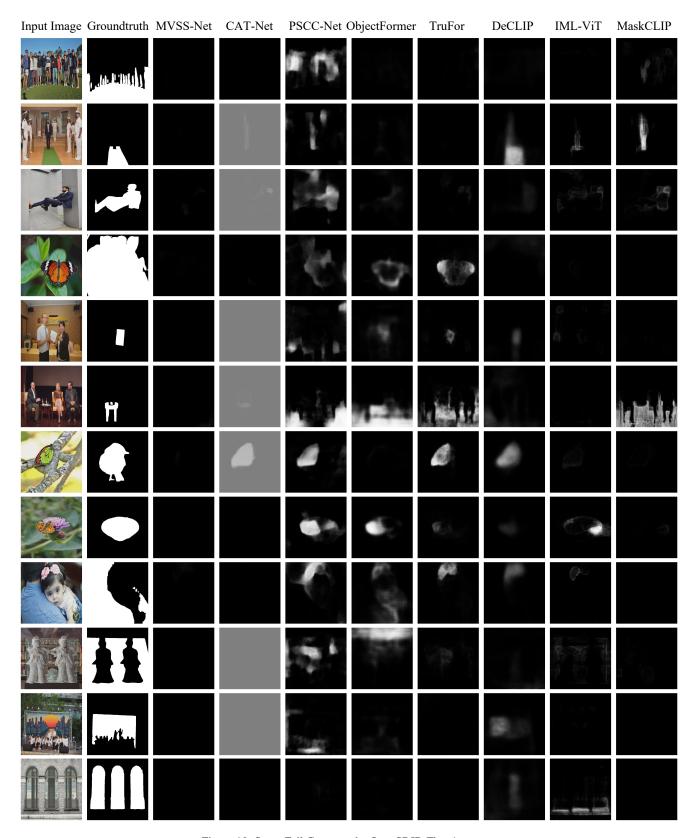


Figure 10. Some Fail Cases on the OpenSDID Flux.1 test set.

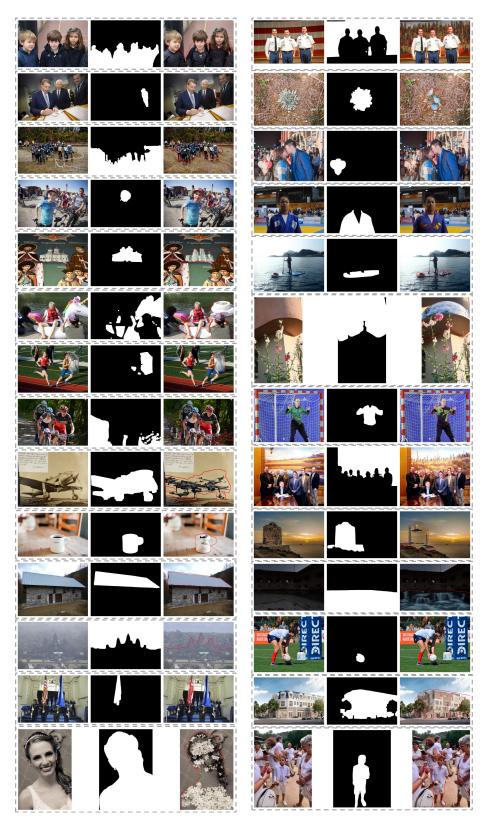


Figure 11. Sample images generated using the Stable Diffusion v1.5 (SD1.5) model.

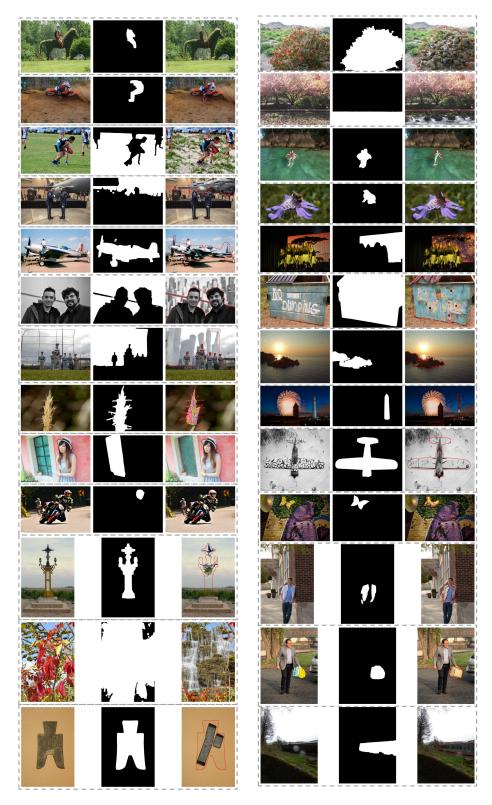


Figure 12. Sample images generated using the Stable Diffusion v2.1 (SD2.1) model.



Figure 13. Sample images generated using the Stable Diffusion v3 model.

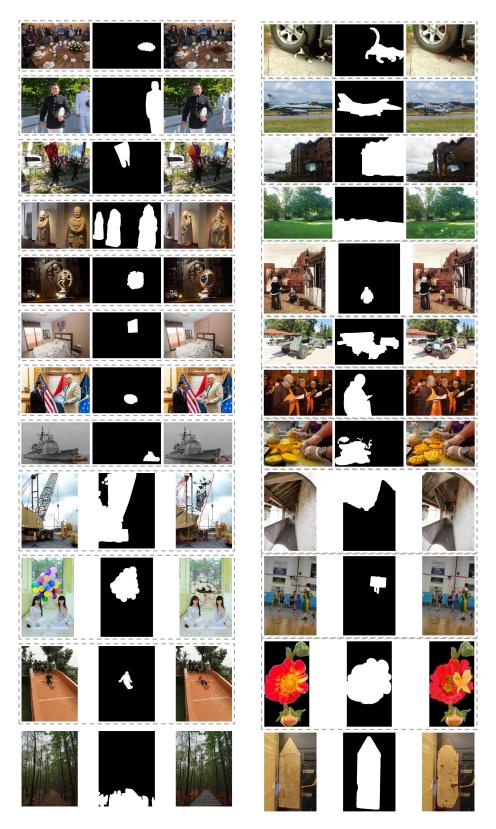


Figure 14. Sample images generated using the Stable Diffusion XL (SDXL) model.

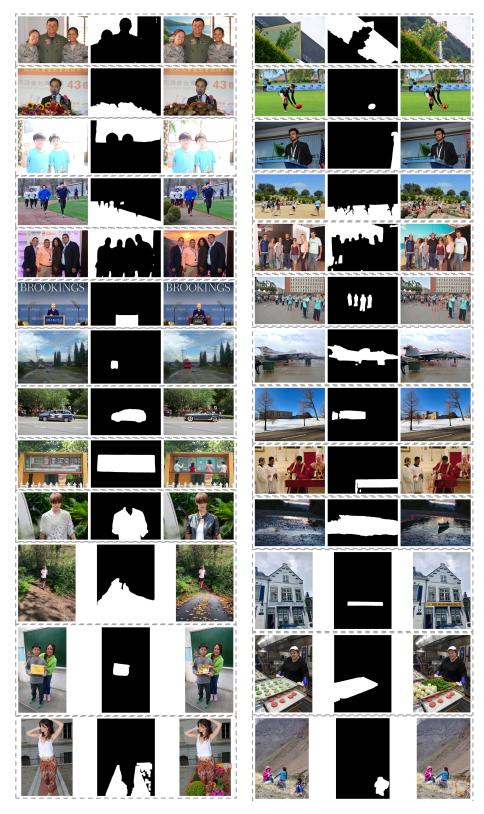


Figure 15. Sample images generated using the Flux.1 model.

- [13] Xiao Guo, Xiaohong Liu, Zhiyuan Ren, Steven Grosz, Iacopo Masi, and Xiaoming Liu. Hierarchical fine-grained image forgery detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3155–3165, 2023. 2, 4
- [14] Shan Jia, Mingzhen Huang, Zhou Zhou, Yan Ju, Jialing Cai, and Siwei Lyu. Autosplice: A text-prompt manipulated image dataset for media forensics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 893–903, 2023. 4
- [15] Liming Jiang, Ren Li, Wayne Wu, Chen Qian, and Chen Change Loy. Deeperforensics-1.0: A large-scale dataset for real-world face forgery detection. In *Proceedings* of the IEEE/CVF conference on computer vision and pattern recognition, pages 2889–2898, 2020. 4
- [16] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Con*ference on Computer Vision, pages 4015–4026, 2023. 1
- [17] Vladimir V Kniaz, Vladimir Knyaz, and Fabio Remondino. The point where reality meets fantasy: Mixed adversarial generators for image splice detection. Advances in neural information processing systems, 32, 2019. 8
- [18] Christos Koutlis and Symeon Papadopoulos. Leveraging representations from intermediate encoder-blocks for synthetic image detection. In *Computer Vision ECCV 2024*, pages 394–411, Cham, 2025. Springer Nature Switzerland. 5
- [19] Myung-Joon Kwon, Seung-Hun Nam, In-Jae Yu, Heung-Kyu Lee, and Changick Kim. Learning jpeg compression artifacts for image manipulation detection and localization. *In*ternational Journal of Computer Vision, 130(8):1875–1895, 2022. 3, 7, 8
- [20] Gant Laborde. Deep nn for nsfw detection. https://github.com/GantMan/nsfw_model, 2019. 3
- [21] Black Forest Labs. Flux.1: A 12 billion parameter rectified flow transformer for text-to-image generation. https://github.com/black-forest-labs/flux, 2024. A 12 billion parameter rectified flow transformer capable of generating images from text descriptions. 2
- [22] Yingxin Lai, Zhiming Luo, and Zitong Yu. Detect any deep-fakes: Segment anything meets face forgery detection and localization. In *Chinese Conference on Biometric Recognition*, pages 180–190. Springer, 2023. 5
- [23] Dong Li, Jiaying Zhu, Xueyang Fu, Xun Guo, Yidi Liu, Gang Yang, Jiawei Liu, and Zheng-Jun Zha. Noise-assisted prompt learning forimage forgery detection and localization. In *Computer Vision – ECCV 2024*, pages 18–36, Cham, 2025. Springer Nature Switzerland. 5
- [24] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023. 1
- [25] Xiaohong Liu, Yaojie Liu, Jun Chen, and Xiaoming Liu. 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, 2022. 4, 7

- [26] Zhengzhe Liu, Xiaojuan Qi, and Philip HS Torr. Global texture enhancement for fake face detection in the wild. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 8060–8069, 2020. 5
- [27] Xiaochen Ma, Bo Du, Xianggen Liu, Ahmed Y Al Hammadi, and Jizhe Zhou. Iml-vit: Image manipulation localization by vision transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024. 5, 6, 7, 8
- [28] Xiaochen Ma, Xuekang Zhu, Lei Su, Bo Du, Zhuohang Jiang, Bingkui Tong, Zeyu Lei, Xinyu Yang, Chi-Man Pun, Jiancheng Lv, et al. Imdl-benco: A comprehensive benchmark and codebase for image manipulation detection & localization. arXiv preprint arXiv:2406.10580, 2024. 3, 6
- [29] Hannes Mareen, Dimitrios Karageorgiou, Glenn Van Wallendael, Peter Lambert, and Symeon Papadopoulos. Tgif: Text-guided inpainting forgery dataset. In 2024 IEEE International Workshop on Information Forensics and Security (WIFS), pages 1–6. IEEE, 2024. 4
- [30] Tian-Tsong Ng, Shih-Fu Chang, and Q Sun. A data set of authentic and spliced image blocks. *Columbia University*, ADVENT Technical Report, 4, 2004. 4
- [31] Adam Novozamsky, Babak Mahdian, and Stanislav Saic. 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, pages 71–80, 2020. 4
- [32] Utkarsh Ojha, Yuheng Li, and Yong Jae Lee. Towards universal fake image detectors that generalize across generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24480–24489, 2023. 5
- [33] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image synthesis. In *The Twelfth International Conference on Learning Representations*, 2024. 2
- [34] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2
- [35] Javier Rando, Daniel Paleka, David Lindner, Lennard Heim, and Florian Tramèr. Red-teaming the stable diffusion safety filter. *arXiv preprint arXiv:2210.04610*, 2022. 3
- [36] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2
- [37] Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. Faceforensics: A large-scale video dataset for forgery detection in human faces. *arXiv preprint arXiv:1803.09179*, 2018. 4
- [38] Christoph Schuhmann and LAION-AI. Clip-based nsfw detector. https://github.com/LAION-AI/CLIP-based-NSFW-Detector, 2022. 3

- [39] Stefan Smeu, Elisabeta Oneata, and Dan Oneata. Declip: Decoding clip representations for deepfake localization. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), 2025. 5, 6, 7, 8
- [40] Chuangchuang Tan, Huan Liu, Yao Zhao, Shikui Wei, Guanghua Gu, Ping Liu, and Yunchao Wei. Rethinking the up-sampling operations in cnn-based generative network for generalizable deepfake detection, 2023. 5
- [41] Chuangchuang Tan, Yao Zhao, Shikui Wei, Guanghua Gu, Ping Liu, and Yunchao Wei. Frequency-aware deepfake detection: Improving generalizability through frequency space learning, 2024. 5
- [42] Dragos-Constantin Tantaru, Elisabeta Oneta, and Dan Oneta. Weakly-supervised deepfake localization in diffusion-generated images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6258–6268, 2024. 4
- [43] Junke Wang, Zuxuan Wu, Jingjing Chen, Xintong Han, Abhinav Shrivastava, Ser-Nam Lim, and Yu-Gang Jiang. Objectformer for image manipulation detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2364–2373, 2022. 5, 7
- [44] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. arXiv preprint arXiv:2409.12191, 2024. 1
- [45] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. Cnn-generated images are surprisingly easy to spot... for now. In *Proceedings of* the IEEE/CVF conference on computer vision and pattern recognition, pages 8695–8704, 2020. 4, 5
- [46] Zijie J Wang, Evan Montoya, David Munechika, Haoyang Yang, Benjamin Hoover, and Duen Horng Chau. Diffusiondb: A large-scale prompt gallery dataset for text-to-image generative models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 893–911, 2023. 2, 4
- [47] Bihan Wen, Ye Zhu, Ramanathan Subramanian, Tian-Tsong Ng, Xuanjing Shen, and Stefan Winkler. Coverage—a novel database for copy-move forgery detection. In 2016 IEEE international conference on image processing (ICIP), pages 161–165. IEEE, 2016. 4
- [48] Yue Wu, Wael AbdAlmageed, and Premkumar Natarajan. Mantra-net: Manipulation tracing network for detection and localization of image forgeries with anomalous features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 7
- [49] Bin Xiao, Haiping Wu, Weijian Xu, Xiyang Dai, Houdong Hu, Yumao Lu, Michael Zeng, Ce Liu, and Lu Yuan. Florence-2: Advancing a unified representation for a variety of vision tasks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4818– 4829, 2024. 1
- [50] Xin Yang, Yuezun Li, and Siwei Lyu. Exposing deep fakes using inconsistent head poses. In ICASSP 2019-2019 IEEE

- International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 8261–8265. IEEE, 2019. 4
- [51] Jizhe Zhou, Xiaochen Ma, Xia Du, Ahmed Y. Alhammadi, and Wentao Feng. Pre-training-free image manipulation localization through non-mutually exclusive contrastive learning. In *Proceedings of the IEEE/CVF International Confer*ence on Computer Vision (ICCV), pages 22346–22356, 2023.
- [52] Mingjian Zhu, Hanting Chen, Qiangyu Yan, Xudong Huang, Guanyu Lin, Wei Li, Zhijun Tu, Hailin Hu, Jie Hu, and Yunhe Wang. Genimage: A million-scale benchmark for detecting ai-generated image. Advances in Neural Information Processing Systems, 36, 2024. 2, 4, 7