

Rescind: Countering Image Misconduct in Biomedical Publications with Vision-Language and State-Space Modeling

Soumyaroop Nandi^{1,2}, Prem Natarajan^{1,2,3}

¹Information Sciences Institute, University of Southern California, Marina Del Rey, CA, USA

²Thomas Lord Department of Computer Science, University of Southern California, Los Angeles, CA, USA

³Capital One, USA

soumyarn@usc.edu, premkumn@usc.edu

Abstract

Scientific image manipulation in biomedical publications poses a growing threat to research integrity and reproducibility. Unlike natural image forensics, biomedical forgery detection is uniquely challenging due to domain-specific artifacts, complex textures, and unstructured figure layouts. We present the first vision-language guided framework for both generating and detecting biomedical image forgeries. By combining diffusion-based synthesis with vision-language prompting, our method enables realistic and semantically controlled manipulations—including duplication, splicing, and region removal—across diverse biomedical modalities. We introduce Rescind, a large-scale benchmark featuring fine-grained annotations and modality-specific splits, and propose Integscan, a structured state-space modeling framework that integrates attention-enhanced visual encoding with prompt-conditioned semantic alignment for precise forgery localization. To ensure semantic fidelity, we incorporate a VLM-based verification loop that filters generated forgeries based on consistency with intended prompts. Extensive experiments on Rescind and existing benchmarks demonstrate that Integscan achieves state-of-the-art performance in both detection and localization, establishing a strong foundation for automated scientific image analysis.

Code — <https://github.com/SoumyaroopNandi/Rescind>

Introduction

Scientific image manipulation has emerged as a significant challenge undermining the integrity of biomedical research. Increasing reports of fraudulent practices—including duplication, splicing, and inappropriate alterations—have led to paper retractions, distorted scientific evidence, and erosion of public trust (Bik, Casadevall, and Fang 2016; Stern et al. 2014). Compared to natural images, biomedical figures—such as microscopy images, gel electrophoresis bands, and histological slides—exhibit complex textures, modality-specific artifacts, and heterogeneous layouts. These challenges are compounded by the rise of generative AI, which enables sophisticated image fabrication (Gu et al. 2022; Kim et al. 2024), making both manual verification and automated detection significantly more difficult.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Forgery synthesis in **BioFors** vs. **Rescind** across biomedical modalities: *macroscopy*, *microscopy*, *FACS*, and *blot/gel*. BioFors uses rectangular insertions lacking realism; Rescind leverages MedSAM (Ma et al. 2024) for irregular mask generation and prompt-guided diffusion for semantically consistent edits. Red = forged; blue = pristine. Eg. Microscopy Prompt: “Duplicate the fluorescent cell cluster from the highlighted region to simulate manipulation.”

Recent advances in image forensics have primarily focused on natural images (Wu, AbdAlmageed, and Natarajan 2019; Kwon et al. 2021; Guillaro et al. 2023; Su et al. 2025), with limited applicability to scientific figures. Moreover, the scarcity of high-quality, publicly available biomedical forgery datasets limits the development and evaluation of robust AI models for scientific integrity applications. Existing biomedical forensics datasets are often small-scale, lack semantic richness and manipulation diversity, and fail to capture realistic domain-specific artifacts (Sabir et al. 2021; Moreira et al. 2022; Mandelli et al. 2022), resulting in poor generalization to real-world scientific image forensics.

Despite recent progress in multimodal foundation models, existing vision-language models such as Biomed-CLIP (Zhang et al. 2024) and LLaVA-Med (Li et al. 2023) are primarily designed for image-text alignment, enabling biomedical image classification, retrieval, and caption generation. However, these models lack the capability to synthesize or manipulate images based on semantic control.

Specifically, while they can accurately describe biomedical images and classify figure modalities, they cannot generate new images exhibiting controlled manipulation types such as duplication or splicing. Furthermore, these models are predominantly trained on radiology-focused datasets, making them more effective at interpreting modalities such as radiographs, MRI, and CT scans, but considerably less robust on microscopy, gel electrophoresis (blot), and other molecular biology figures commonly used in biomedical publications. This significant gap limits their applicability in generating realistic forgery datasets or detecting scientific image manipulations within non-radiological biomedical contexts.

To address these challenges, we propose a VLM-guided forgery generation and detection framework for biomedical publications. We introduce **Rescind**, a large-scale benchmark dataset of synthetically manipulated biomedical images spanning diverse modalities, including microscopy, blot/gel, macroscopy, and flow cytometry (FACS). Rescind is constructed using a VLM-driven semantic control mechanism combined with a diffusion-based generative pipeline to produce realistic forgery types, including duplication, splicing, and region removal. Furthermore, we incorporate a VLM-based automatic verification loop to ensure semantic consistency and quality control of the generated forgeries, addressing a critical gap in existing data curation approaches for scientific image integrity analysis.

In addition to dataset generation, we introduce **Integscan**, a vision-language architecture for fine-grained biomedical image forgery detection. Integscan integrates structured state-space modeling with lightweight attention mechanisms to encode spatial features, while leveraging vision-language prompts for semantic alignment. Its modular design—including IntegSSM blocks, prompt-conditioned enhancement via IntegBoost, and structured fusion through the Integration module—enables accurate localization of manipulated regions across diverse modalities. To ensure semantic realism in forged samples, we fine-tune diffusion models for region-aware inpainting and incorporate a VLM-based verification loop for filtering. Extensive experiments on Rescind and existing biomedical benchmarks demonstrate that Integscan achieves state-of-the-art performance in both pixel and image-level forgery detection and localization. Our key contributions are:

- We introduce **Rescind**, a 600K-scale benchmark for biomedical image forgery detection, comprising three synthetic variants—**Rescind-I** (image-processed), **Rescind-G** (generative), and **Rescind-V** (VLM-guided)—spanning four manipulation types across multiple biomedical modalities.
- We present the first VLM-guided forgery synthesis pipeline that combines vision-language prompting, generative diffusion, and a post-hoc verification loop using LLaVA-Med to produce semantically aligned and visually realistic biomedical forgeries.
- We propose **Integscan**, a state-space model with attention-enhanced visual encoding and prompt-conditioned semantic fusion, achieving state-of-the-art localization across all manipulation tasks and modalities.

Related Work: Biomedical Misconduct

Limitations of Existing Datasets. Image manipulation is a growing concern in biomedical publishing, with up to 3.8% of papers estimated to contain tampered images (Bik, Casadevall, and Fang 2016). While synthetic datasets like BioFors (Sabir et al. 2021), Mandelli et al. (Mandelli et al. 2022), and M3D-Synth (Zingarini et al. 2024) have supported detection efforts, they are constrained by modality (e.g., Western blots, CT scans), lack pixel-level annotations, or simulate manipulations using simplistic heuristics. To overcome these limitations, we introduce **Rescind**, a general-purpose biomedical forgery dataset with fine-grained masks spanning microscopy, blot/gel, macroscopy, and FACS. Rescind includes semantically meaningful, irregular manipulations grounded in real figures, making it well-suited for training and evaluation.

Towards General-Purpose Forgery Detection. Despite advances in biomedical segmentation (Shamshad et al. 2023; Heidari et al. 2023; Wu et al. 2022) and diagnosis (Esteva et al. 2021), biomedical forensics remains underexplored due to limited data and task diversity. Early duplication detection methods (Koppers, Wormer, and Ickstadt 2017; Bucci 2018; Acuna, Brookes, and Kording 2018) based on hand-crafted features perform poorly in domains with subtle textures and structural repetition. Deep models (Wu, AbdAlmageed, and Natarajan 2019; Guillaro et al. 2023; Su et al. 2025) have improved performance but often rely on splicing assumptions and fail under complex internal forgeries. While BioFors (Sabir et al. 2021) introduced EDD, IDD, and CSTD benchmarks, most models—e.g., MONet (Sabir et al. 2022)—focus on EDD alone. To this end, we propose **IntegScan**, a vision-language-guided state-space model that jointly localizes source and manipulated regions across all three BioFors tasks, offering improved generalization to diverse biomedical forgery types.

VLMs for Biomedical Images. VLMs like CLIP (Radford et al. 2021) and its biomedical variants—BiomedCLIP (Zhang et al. 2024), OpenBioMed (Luo et al. 2023), and LLaVA-Med (Li et al. 2023)—excel in classification, retrieval, and captioning. However, they are not designed to detect manipulation or exposed to forged data. We fine-tune VLMs on Rescind with manipulation-aware prompts, enabling semantically conditioned forgery detection and enhancing both interpretability and realism.

State Space Models. Structured State Space Models (SSMs), such as S4 (Gu, Goel, and Ré 2021), Mamba (Gu and Dao 2023), and Mamba-2 (Dao and Gu 2024), model long-range dependencies efficiently and are increasingly viewed as generalizations of Transformers (Han et al. 2024). Vision extensions like ViMamba (Liu et al. 2024) have shown promise in segmentation, but remain generic. We introduce the first SSM-based model for biomedical image forensics, leveraging structured recurrence and attention to achieve high accuracy on complex forgery types.

Rescind Dataset Construction

We introduce Rescind, a large-scale dataset of semantically controlled and visually realistic forgeries in biomedical im-

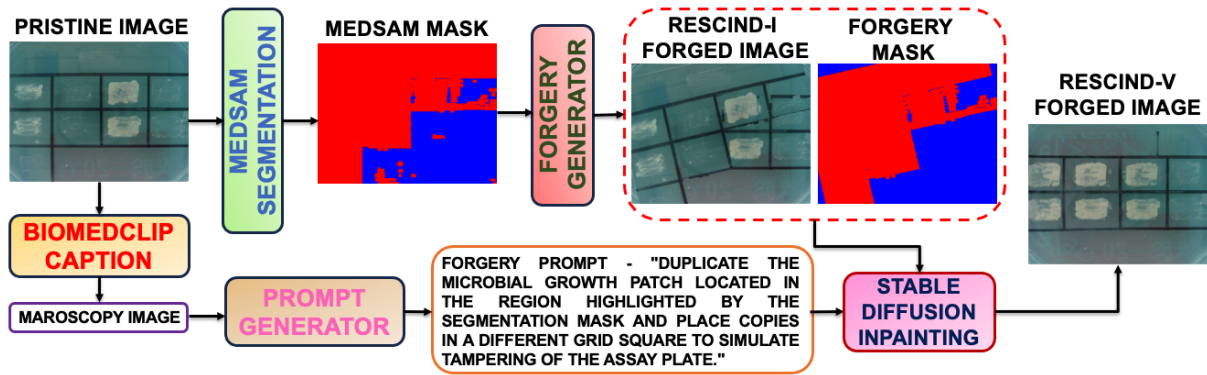


Figure 2: **Rescind-V Curation Pipeline.** Given a *pristine biomedical image*, BiomedCLIP (Zhang et al. 2024) generates a modality-aware caption (e.g., *macroscopy image*). MedSAM (Ma et al. 2024) produces a binary segmentation mask that identifies regions of interest. A **Forgery Generator** creates a Rescind-I forged image and corresponding forgery mask by duplicating or manipulating content within the masked region. A **Prompt Generator** formulates a semantic forgery prompt based on the predicted modality, explicitly referencing the MedSAM-highlighted region. Finally, a Stable Diffusion inpainting model (Rombach et al. 2022) synthesizes the **Rescind-V forged image** using the *forgery mask* and *forgery prompt* as conditional inputs.

ages. Rescind is curated through a multi-stage pipeline that integrates vision-language understanding, region segmentation, and guided image generation to produce high-fidelity forgeries while maintaining biological plausibility. Figure 2 provides an overview of the complete generation pipeline.

Pristine Image Collection

We begin by curating a diverse set of authentic biomedical figures from publicly available and ethically approved datasets. Specifically, 1,031 Portable Document Format (PDF) files containing raw annotations of suspicious scientific figures were obtained from (Bik, Casadevall, and Fang 2016), of which 335 contained at least one manipulated figure. Following the approach in (Sabir et al. 2021), we extracted 17,269 manipulated images (**Rescind-T**) and 30,536 unmanipulated pristine images (**Rescind-P**) from these figures. The images span a range of biomedical modalities and are categorized into four major classes based on semantics/visual appearance as illustrated in first column of Figure 1:

Microscopy comprises images captured using light or electron microscopes, typically depicting stained cells and tissues. These images vary widely in color and structure due to differences in biological source and staining techniques, contributing to their visual heterogeneity.

Blots/Gels include western, northern, and southern blot images and electrophoresis gels used for analyzing proteins, RNA, and DNA. Despite targeting different biomolecules, they exhibit similar visual formats such as bands and lanes, often indistinguishable without textual labels.

Flow Cytometry (FACS) images resemble scatter plots, showing cell or particle distributions captured via fluorescence-based experiments. Though visually simple, they encode biologically meaningful cell population data.

Macroscopy includes biomedical images captured without magnification, such as gross anatomical views, surgical fields, or laboratory setups, and exhibits high variability in structure and context.

These modality-aware, richly annotated pristine images form the foundation for generating both authentic and synthetic samples in the Rescind dataset.

Synthetic Forgery Generation

Building on the modality-aware pristine figures in **Rescind-P**, we construct three categories of synthetically manipulated images, each simulating different manipulation strategies commonly encountered in biomedical research. These subsets—**Rescind-I**, **Rescind-G**, and **Rescind-V**—enable diverse and controllable image forgery types ranging from low-level transformations to semantically guided edits.

Rescind-I (Image-processed forgeries) contains forged images created by applying image-level transformations to segmented regions extracted from Rescind-P using MedSAM (Ma et al. 2024). These transformations include scaling, rotation, shifting, and duplication of image patches, which are either repositioned within the same image (Internal Duplication Detection - IDD) or blended into new host images using MGMatting (Yu et al. 2021) (External Duplication Detection - EDD). Unlike prior datasets with rectangular masks (e.g., BioFors (Sabir et al. 2021)), Rescind-I employs irregular masks that better reflect natural biological structures, improving realism and localization difficulty.

Rescind-G (Generative forgeries) builds upon Rescind-P and Rescind-I by introducing forgery types generated via image-to-image generative models. Specifically, we apply Stable Diffusion inpainting (Rombach et al. 2022), Pix2Pix GAN (Isola et al. 2017), and CycleGAN (Zhu et al. 2017) to MedSAM-derived masks. These methods generate realistic content in place of or in addition to existing regions, simulating removal or substitution forgeries. The generated patches are applied either to the original region defined by MedSAM or to new locations introduced in Rescind-I, allowing spatial and contextual variability.

Rescind-V (VLM-guided semantic forgeries) introduces semantic control into forgery synthesis by leveraging vision-

Subset	Blot	Micro	Macro	FACS	Total
Rescind-P (Train)	19,105	10,458	555	418	30,536
Rescind-T (Test)	8,335	7,652	639	643	17,269
Rescind-I (EDD)	40,002	25,002	2,502	2,502	70,008
Rescind-I (IDD)	9,999	6,501	1,251	1,251	19,002
Rescind-I (CSTD)	9,999	6,501	1,251	1,251	19,002
Rescind-I (Removal)	40,002	25,002	2,502	2,502	70,008
Rescind-G (EDD)	40,002	25,002	2,502	2,502	70,008
Rescind-G (IDD)	9,999	6,501	1,251	1,251	19,002
Rescind-G (CSTD)	9,999	6,501	1,251	1,251	19,002
Rescind-G (Removal)	40,002	25,002	2,502	2,502	70,008
Rescind-V (EDD)	40,002	25,002	2,502	2,502	70,008
Rescind-V (IDD)	9,999	6,501	1,251	1,251	19,002
Rescind-V (CSTD)	9,999	6,501	1,251	1,251	19,002
Rescind-V (Removal)	40,002	25,002	2,502	2,502	70,008
Rescind-Total	327,446	207,128	23,712	23,579	581,865

Table 1: Rescind Summary: pristine(P), real test(T), and synthetic subsets—Image-processed(I), Generative(G), VLM-guided(V); Task types: EDD (external duplication), IDD (internal duplication), CSTD (cut/sharp transition), Removal.

language models (VLMs). Each image in Rescind-P is first captioned using BiomedCLIP (Zhang et al. 2024) to extract modality-aware descriptions. MedSAM (Ma et al. 2024) is then applied to segment regions of interest, which are paired with curated natural language prompts describing plausible manipulations (e.g., “duplicate the protein band gel lane identified by the mask and replicate it in a new position to fabricate consistent results across experiments” or “Erase the mitotic cell cluster from the highlighted area to simulate data suppression in the microscopy image.”).

These prompts are generated by a Prompt Generator, which conditions on the BiomedCLIP-derived caption and incorporates manipulation patterns inspired by the forensic taxonomy of notes outlined in (Bik, Casadevall, and Fang 2016) and real-world examples from retracted publications. To guide prompt construction, we also studied the notes and annotations accompanying detected forgeries in retracted papers, enabling the generation of realistic semantic prompts that reflect tampering intent from a perpetrator’s perspective.

A Forgery Generator then uses the MedSAM-produced mask to create an initial Rescind-I forged image and corresponding forgery mask. While Rescind-I introduces structural changes, these forgeries are not always semantically coherent. To overcome this, Rescind-V employs Stable Diffusion inpainting (Rombach et al. 2022), conditioned on the pristine image, the forgery mask (which guides the inpainting region), and the semantic prompt, to synthesize realistic and contextually consistent forgeries. This subset emphasizes language-aligned, modality-consistent manipulation and supports the development and evaluation of VLM-sensitive forgery detection models.

Together, Rescind-I, Rescind-G, and Rescind-V form a unified benchmark for image integrity, spanning low-level duplication, generative manipulation, and semantically guided forgery synthesis, as summarized in Table 1.

Proposed Method

Integscan is a VLM architecture for fine-grained biomedical image forgery detection. It integrates structured spatial encoding and prompt-conditioned semantic alignment to localize manipulated regions across diverse modalities.

Integscan Pipeline Overview

As illustrated in Figure 3, given a forged Rescind-V image $V \in \mathbb{R}^{b \times c \times h \times w}$, its corresponding forgery mask, and an associated prompt, **Integscan** predicts a binary manipulation mask $\hat{O} \in \mathbb{R}^{b \times c \times h \times w}$, where b , c , h , and w denote batch size, channels, height, and width, respectively. A structured state-space encoder first extracts visual features from V , which are processed by stacked **IntegSSM** blocks to generate flattened spatial tokens $X_v \in \mathbb{R}^{b \times N \times D}$, where $N = H_f \times W_f$ is the number of downsampled spatial tokens and D is the embedding dimension.

To incorporate modality-specific semantics, we leverage BiomedCLIP (Zhang et al. 2024) to retrieve relevant prompts and embed them alongside visual features. These embeddings are passed to the **IntegBoost** module, which performs semantic enhancement using structured attention and state-space reasoning to produce enriched prompt-aware features $X_t \in \mathbb{R}^{b \times N \times 2D}$. The enhanced text features are fused with the visual tokens using the **Integration** module, yielding a unified representation $X_{\text{int}} \in \mathbb{R}^{b \times N \times 2D}$. Finally, X_{int} is reshaped and passed through a lightweight convolutional decoder comprising a ReLU-activated 3×3 layer followed by a 1×1 projection to generate the dense manipulation mask. We detail the core visual encoding and attention components of Integscan below:

IntegSSM Block: Structured Visual Encoding

The **IntegSSM** block forms the backbone of the visual encoder in Integscan. It sequentially integrates convolutional and attention-based operations to encode spatial structure at multiple levels of granularity. As shown in Figure 4(a), the input tokens first pass through a CONV layer to extract local spatial features. This is followed by two successive LAYERNORM layers, interleaved with additional CONV operations to refine the representation. A SILU activation introduces non-linearity before the tokens are passed into the INTEG_ATT_N submodule (detailed in Section), which models long-range spatial dependencies using linear attention with rotary positional encoding. The output is then projected via a LINEAR layer and passed through another CONV \rightarrow LAYERNORM \rightarrow MLP sequence to further mix token interactions across space and channels. The final output is produced by a concluding LINEAR projection. This block design enables efficient extraction of hierarchical features while preserving both local and global contextual cues.

Integ Attn: Structured Linear Attention via SSM

At the heart of the IntegSSM block lies the `IntegAttn` submodule, which implements structured linear attention based on discretized State Space Models (SSMs). This formulation allows efficient and scalable modeling of long-range spatial dependencies through learned recurrence and

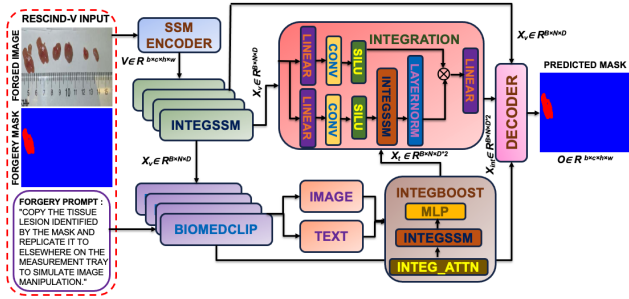


Figure 3: Integscan Model Architecture Overview

positional priors. The continuous-time SSM is defined as:

$$h'(t) = \mathbf{A}h(t) + \mathbf{B}x(t), \quad y(t) = \mathbf{C}h(t) \quad (1)$$

To integrate this into deep networks, we apply zero-order hold (ZOH) discretization (Gu, Goel, and Ré 2021), resulting in:

$$h_k = \bar{\mathbf{A}}h_{k-1} + \bar{\mathbf{B}}x_k, \quad y_k = \bar{\mathbf{C}}h_k + \bar{\mathbf{D}}x_k \quad (2)$$

where $\bar{\mathbf{A}}, \bar{\mathbf{B}}, \bar{\mathbf{C}}$ are learned discrete parameters and $\bar{\mathbf{D}}$ provides a residual shortcut.

In practice, `IntegAttn` approximates these dynamics through a series of differentiable steps. The input is projected into query and key vectors, activated with ELU and offset by a bias term to ensure positive gating. Rotary positional encoding (RoPE) is applied to both q and k to encode spatial structure in a recurrent form. Token-wise recurrence is simulated by computing attention weights through interactions between q and a compressed average of k . These weights are used to aggregate value tokens, which are further enhanced with a depthwise convolutional layer—analogueous to the residual $\bar{\mathbf{D}}x_k$ term in Equation 2. By combining structured recurrence, positional priors, and local spatial enhancement, `IntegAttn` provides a lightweight yet expressive mechanism for structured attention within Integscan.

IntegBoost: Prompt Enhance via Visual Context

The `IntegBoost` module, shown in Figure 3, enhances prompt embeddings by conditioning them on visual priors extracted from the forged image. Unlike previous methods that treat text as static queries, `IntegBoost` dynamically refines the prompts by modeling their interaction with image features. It consists of a lightweight attention module, a structured `IntegSSM` block, and an MLP, each equipped with residual connections. The attention layer first captures contextual alignment between the prompt and visual tokens, while the `IntegSSM` block further integrates spatial priors into the prompt representation. This hybrid design improves semantic grounding, allowing the prompt to effectively guide manipulation localization across diverse biomedical modalities.

Integration Module: Visual–Prompt Feat. Fusion

The `Integration` module, shown in Figure 3, fuses visual features from `IntegSSM` with prompt-enhanced representations from `IntegBoost` to produce a unified token sequence




Figure 4: (a) IntegSSM with state-space modeling; (b) INTEG_ATT with rotary encoding, channel-aware attention.

for manipulation localization. Rather than relying on cross-attention over concatenated tokens—which scales poorly with sequence length—the module adopts a structured design based on linear projections, convolutional refinement, and state-space modeling. Both inputs are first projected into a shared embedding space and reshaped into spatial grids. Each stream is individually processed by `SILU`-activated convolution layers to capture local structure, then element-wise summed to encourage early feature interaction. The result is passed through an `IntegSSM` block, which models spatial dependencies across the fused representation via structured recurrence. A final normalization and linear projection produce the fused embedding $X_{\text{int}} \in \mathbb{R}^{b \times N \times 2D}$ for downstream decoding. Inspired by adapter-based learning (Houlsby et al. 2019), this approach allows visual and prompt features to interact meaningfully without fully fine-tuning large pretrained encoders. The state-space fusion captures both long-range token relationships and local visual cues while maintaining computational efficiency. This makes `Integration` a lightweight yet expressive module for prompt-guided segmentation.

Experimental Evaluation

Rescind Quality Assessment & VLM Validation

To ensure the realism and utility of the `Rescind` dataset, we conduct both perceptual and semantic quality evaluations of the generated forgeries.

Visual Quality Evaluation. We assess image fidelity using both no-reference and full-reference image quality metrics. For forged images alone, we employ no-reference metrics such as NIQE (Ying et al. 2020) and PIQE (Wang, Chan, and Loy 2023) to evaluate perceptual naturalness in the absence of ground truth. For comparing forged images against their pristine counterparts, we use full-reference metrics including SSIM (Wang et al. 2004) and LPIPS (Zhang et al. 2018) to measure structural similarity and perceptual distance. To further eliminate trivially unrealistic samples, we train a binary image classifier to detect synthetic forgeries, filtering out generated samples with high fake-confidence scores to maintain dataset integrity. Table 2 reports the results of these quality assessments on the synthetic subsets of the `Rescind` dataset. Among all subsets, **Rescind-V** achieves the best scores across all metrics—lowest NIQE and PIQE, highest SSIM, and lowest LPIPS—indicating that vision-language

Metric	Type	Rescind-I	Rescind-G	Rescind-V
NIQE (↓)	No-Ref	4.23	3.79	3.31
PIQE (↓)	No-Ref	38.7	30.5	25.9
SSIM (↑)	Full-Ref	0.827	0.871	0.894
LPIPS (↓)	Full-Ref	0.204	0.143	0.116

Table 2: Image Quality Assessment (IQA) on Rescind subsets. ↑ indicates higher is better; ↓ indicates lower is better.

model (VLM)-guided forgeries are the most visually plausible and semantically coherent. This underscores the importance of incorporating modality-aware semantic control in generating high-fidelity biomedical forgeries.

Semantic Consistency Verification. Beyond visual realism, we ensure that each manipulation aligns with its intended semantic objective. To this end, we employ LLaVA-Med (Li et al. 2023) in a post-hoc verification loop. Each generated image is captioned by the VLM using a modality-specific manipulation prompt (e.g., “image with duplicated blot band”), and the output is compared against the expected semantic intent. Images falling below a predefined alignment threshold are discarded. The remaining forgeries are annotated with a *Semantic Consistency Score*, which quantifies alignment quality for downstream analysis. Figure 5 shows score distributions across forgery types, with duplication exhibiting the highest semantic fidelity.

Integscan Results and Analysis

Integscan achieves state-of-the-art performance on the Rescind-T test set, containing real-world duplication forgeries from retracted biomedical publications. As shown in Tables 3-4, it consistently outperforms prior methods across all biomedical modalities and tasks, including both External and Internal Duplication Detection (EDD and IDD). Unlike MONet (Sabir et al. 2022), which is limited to EDD, Integscan delivers robust performance across all categories.

Traditional keypoint-based detectors (e.g., SIFT (Lowe 2004), ORB (Rublee et al. 2011)) and handcrafted feature methods (e.g., DCT (Fridrich et al. 2003), Zernike (Ryu, Lee, and Lee 2010)) perform poorly on biomedical images due to structural repetition. While dense matching methods like DF-ZM and DF-PCT (Cozzolino, Poggi, and Verdoliva 2015) outperform earlier techniques, their reliance on low-level features limits their generalization. Deep models such as ManTra-Net (Wu, AbdAlmageed, and Natarajan 2019), TruFor (Guillaro et al. 2023), and MONet struggle to handle duplication-based tampering, particularly when confined to target-only prediction.

Integscan uniquely localizes both source and target regions in duplication forgeries, including challenging internal duplication scenarios. While competing models often misclassify biomedical textures or detect only superficial manipulation artifacts, Integscan accurately identifies contextually aligned forgeries with high spatial precision. This is achieved by combining structured attention with state-space modeling and prompt-guided semantics, enabling robust visual reasoning across diverse biomedical modalities.




Figure 5: Semantic Consistency Score across forgery types, using LLaVA-Med on Rescind-I, Rescind-G, Rescind-V.

Method	Internal Duplication Detection (IDD)							
	Microscopy Image Pixel		Blot/Gel Image Pixel		Macroscopy Image Pixel		Combined Image Pixel	
DF-ZM	<u>0.764</u>	0.197	0.515	0.449	0.573	0.478	0.564	0.353
DF-PCT	<u>0.764</u>	0.202	0.503	0.466	0.712	0.487	0.569	0.364
DF-FMT	0.638	0.167	0.480	0.400	0.495	0.458	0.509	0.316
BusterNet	0.183	0.178	0.226	0.076	0.021	0.106	0.269	0.107
ManTra	0.316	0.194	0.317	0.094	0.272	0.262	0.335	0.183
TruFor	0.321	0.197	0.324	0.099	0.245	0.252	0.336	0.197
Integscan	0.906	0.604	0.768	0.694	0.925	0.764	0.793	0.614

Table 3: MCC scores for IDD on Rescind-T test-set for both Image and Pixel levels; best **bold**, second-best underlined

Ablation Studies

Semantic Prompt Importance. We ablate the vision-language component by generating Rescind-V samples without BiomedCLIP-guided prompts. Instead, we apply random or generic prompts (e.g., “remove a region”). Detection models trained on prompt-less samples perform worse on semantically consistent manipulations, confirming the critical role of prompt conditioning for generating realistic, language-aligned forgeries.

Modality-Specific Generalization. We perform a leave-one-modality-out experiment where models are trained on three modalities and evaluated on the held-out fourth. This reveals cross-modality generalization gaps and highlights the need for modality-aware modeling. For example, models trained without FACS perform poorly on FACS test samples, indicating that visually simple modalities require explicit inclusion in training.

Integscan Ablation. To understand the contribution of each component in the **Integscan** architecture, we conduct systematic ablations on the Rescind-T EDD test set (Table 6). Removing the **IntegBoost** module, which enhances text embeddings through image-conditioned attention, leads to a noticeable drop in F1 score and IoU, highlighting the importance of modality-aware prompt refinement. Excluding the **Integration** fuser, which bridges VLM and VFM features, reduces the model’s ability to semantically align predictions, resulting in degraded performance. When the **IntegSSM** visual backbone is replaced with a standard convolutional encoder, the model’s spatial reasoning capacity is significantly reduced, confirming the advantage of long-range attention

Method	External Duplication Detection (EDD)									
	Micro		Blot/Gel		Macro		FACS		Comb.	
	Img	Px	Img	Px	Img	Px	Img	Px	Img	Px
SIFT	0.180	0.146	0.113	0.148	0.130	0.194	0.110	0.073	0.142	0.132
ORB	0.319	0.342	0.087	0.127	0.126	0.226	0.269	0.187	0.207	0.252
BRIEF	0.275	0.277	0.058	0.102	0.135	0.169	0.244	0.188	0.180	0.202
DF-ZM	0.422	0.425	0.161	0.192	0.285	0.256	0.540	0.504	0.278	0.324
DMVN	0.342	0.242	0.430	0.261	0.238	0.185	0.282	0.164	0.310	0.244
ManTra	0.347	0.244	0.449	0.287	0.275	0.202	0.337	0.186	0.351	0.231
TruFor	0.356	0.261	0.468	0.299	0.214	0.357	0.349	0.189	0.409	0.317
MONet	0.435	0.398	0.520	0.507	0.262	0.221	0.356	0.313	0.438	0.410
Ours	0.821	0.556	0.755	0.662	0.818	0.654	0.726	0.507	0.776	0.604

Table 4: MCC scores for EDD on Rescind-T test-set for both Image and Pixel levels; best **bold**, second-best underlined

Training Configuration	Precision	Recall	F1	IoU
w/o BiomedCLIP Prompts	75.92	74.18	75.04	61.20
w/o Microscopy (test Mic.)	72.35	68.49	70.36	55.81
w/o Blots/Gels (test Blot/Gel)	73.12	71.18	72.14	58.26
w/o FACS (test FACS)	66.88	60.57	63.40	49.75
w/o Macroscopy (test Macro)	70.15	67.03	68.55	54.39
Full Dataset	79.82	78.05	78.92	63.48

Table 5: Dataset ablation on Rescind-V using Integscan.

via structured state-space modeling. Finally, bypassing the prompt-conditioning process and using a generic prompt across all images leads to a marked decline in localization performance, verifying the necessity of semantic prompts tailored to the image modality. These results collectively validate the design choices of Integscan in achieving precise and robust forgery localization.

Robustness Analysis

To assess the robustness of forgery detection models under real-world degradations, we evaluate all baselines on the **Rescind-V** dataset across six common perturbations—*JPEG compression*, *Gaussian blurring*, *Gaussian noise*, *brightness shift*, *contrast adjustment*, and *cropping*—following the protocol in (Wu, Abd-Almageed, and Natarajan 2018; Wu, AbdAlmageed, and Natarajan 2019). Figure 6 presents the pixel-level F1 scores for five representative detection methods under increasing perturbation severity. **Integscan** demonstrates consistently superior performance across all attack types, exhibiting graceful degradation even under heavy distortions. In contrast, existing methods such as SparseViT, MVSS, and TruFor show significant performance drops under geometric and low-level corruptions. These results confirm that Integscan is not only accurate

Model Variant	Precision	Recall	F1 Score	IoU
w/o IntegBoost	94.62	92.88	93.74	88.86
w/o Integration Module	92.19	90.72	91.45	85.33
w/o Prompt Conditioning	91.74	89.31	90.51	83.94
IntegSSM → CNN	90.28	88.75	89.51	81.80
Integscan (Full)	97.83	96.94	97.38	94.96

Table 6: Ablation study on the Rescind-T EDD images.




Figure 6: Robustness under attacks on Rescind-V

but also highly robust to post-manipulation operations commonly found in scientific image tampering.

Limitations and Future Work

While Integscan achieves strong performance across diverse biomedical forgery detection tasks, some limitations remain. First, reliance on modality-specific prompts inferred from a pretrained VLM can lead to suboptimal alignment if the modality is misclassified—particularly for ambiguous or low-quality inputs. Second, although our synthetic dataset captures a broad range of manipulation types, it may not fully represent compositional or multi-step forgeries seen in real-world fraud, such as sequential tampering across figure panels or coordination between image and caption. Lastly, the use of a shared encoder across forgery types may constrain the model’s ability to specialize in task-specific cues. Future work will explore retrieval- or reinforcement-based prompt refinement, validation across a broader set of retracted publications, and reviewer-in-the-loop feedback to support deployment in high-stakes forensic pipelines.

Conclusion

We introduce a vision-language guided framework for generating and detecting semantic forgeries in biomedical images. By combining diffusion-based synthesis, vision-language prompting, and structured state-space modeling, we construct **Rescind**—the first large-scale benchmark enabling semantically controlled forgeries across diverse biomedical modalities. To ensure semantic fidelity, we add a vision-language verification loop during generation. We also propose **Integscan**, a model combining attention-enhanced encoding with prompt-conditioned fusion, achieving state-of-the-art performance across manipulation detection tasks. Code and dataset will be released to support future research in scientific integrity, establishing a scalable foundation for trustworthy and reproducible AI in biomedical forensics.

Ethics Statement

This work aims to strengthen scientific integrity by improving detection of manipulated biomedical images. All data in Rescind are sourced from publicly available datasets and retracted papers, with no private or patient-identifiable information, and used in accordance with original licenses. To mitigate dual-use risks, all generated images are metadata-tagged, and forgery generation tools will be released only under controlled access for verified research use. We follow responsible stewardship and NIST AI 800-1 guidelines, including license restrictions and rate limits. Malicious actors could theoretically study the benchmark to identify detector weaknesses; however, we believe the benefit of enabling defensive research outweighs this risk. The societal goal is to help journals, integrity offices, and the research community identify fraudulent figures and improve trust in publications. If any party of concern objects to specific images, we will remove them promptly. We acknowledge that synthetic forgeries may not capture the full diversity of real-world misconduct and encourage future work on broader validation.

References

- Acuna, D. E.; Brookes, P. S.; and Kording, K. P. 2018. Bioscience-scale automated detection of figure element reuse. *BioRxiv*, 269415.
- Bik, E. M.; Casadevall, A.; and Fang, F. C. 2016. The prevalence of inappropriate image duplication in biomedical research publications. *MBio*, 7(3): 10–1128.
- Bucci, E. M. 2018. Automatic detection of image manipulations in the biomedical literature. *Cell death & disease*, 9(3): 400.
- Cozzolino, D.; Poggi, G.; and Verdoliva, L. 2015. Efficient dense-field copy-move forgery detection. *IEEE Transactions on Information Forensics and Security*, 10(11): 2284–2297.
- Dao, T.; and Gu, A. 2024. Transformers are SSMs: Generalized Models and Efficient Algorithms Through Structured State Space Duality. In *International Conference on Machine Learning (ICML)*.
- Esteva, A.; Chou, K.; Yeung, S.; Naik, N.; Madani, A.; Motlaghi, A.; Liu, Y.; Topol, E.; Dean, J.; and Socher, R. 2021. Deep learning-enabled medical computer vision. *NPJ digital medicine*, 4(1): 5.
- Fridrich, J.; Soukal, D.; Lukas, J.; et al. 2003. Detection of copy-move forgery in digital images. In *Proceedings of digital forensic research workshop*, volume 3, 652–63. Cleveland, OH.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Gu, A.; Goel, K.; and Ré, C. 2021. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*.
- Gu, J.; Wang, X.; Li, C.; Zhao, J.; Fu, W.; Liang, G.; and Qiu, J. 2022. AI-enabled image fraud in scientific publications. *Patterns*, 3(7).
- Guillaro, F.; Cozzolino, D.; Sud, A.; Dufour, N.; and Verdoliva, L. 2023. Trufor: Leveraging all-round clues for trustworthy image forgery detection and localization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 20606–20615.
- Han, D.; Wang, Z.; Xia, Z.; Han, Y.; Pu, Y.; Ge, C.; Song, J.; Song, S.; Zheng, B.; and Huang, G. 2024. Demystify Mamba in Vision: A Linear Attention Perspective. In *NeurIPS*.
- Heidari, M.; Kazerouni, A.; Soltany, M.; Azad, R.; Aghdam, E. K.; Cohen-Adad, J.; and Merhof, D. 2023. Hi-former: Hierarchical multi-scale representations using transformers for medical image segmentation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 6202–6212.
- Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; De Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly, S. 2019. Parameter-efficient transfer learning for NLP. In *International Conference on Machine Learning (ICML)*, 2790–2799. PMLR.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1125–1134.
- Kim, J. J.; Um, R. S.; Lee, J. W.; and Ajilore, O. 2024. Generative AI can fabricate advanced scientific visualizations: Ethical implications and strategic mitigation framework. *AI and Ethics*, 1–13.
- Koppers, L.; Wormer, H.; and Ickstadt, K. 2017. Towards a systematic screening tool for quality assurance and semiautomatic fraud detection for images in the life sciences. *Science and engineering ethics*, 23(4): 1113–1128.
- Kwon, M.-J.; Yu, I.-J.; Nam, S.-H.; and Lee, H.-K. 2021. CAT-Net: Compression artifact tracing network for detection and localization of image splicing. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 375–384.
- Li, C.; Wong, C.; Zhang, S.; Usuyama, N.; Liu, H.; Yang, J.; Naumann, T.; Poon, H.; and Gao, J. 2023. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *Advances in Neural Information Processing Systems*, 36: 28541–28564.
- Liu, Y.; Tian, Y.; Zhao, Y.; Yu, H.; Xie, L.; Wang, Y.; Ye, Q.; and Liu, Y. 2024. VMamba: Visual State Space Model. *arXiv preprint arXiv:2401.10166*.
- Lowe, D. G. 2004. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60: 91–110.
- Luo, Y.; Yang, K.; Fan, S.; Hong, M.; Zhao, S.; Chen, X.; Nie, Z.; Luo, W.; Xie, A.; Liu, X. Y.; Zhang, J.; Wu, Y.; and Nie, Z. 2023. Code of OpenBioMed. <https://github.com/Pharmolix/OpenBioMed.git>.
- Ma, J.; He, Y.; Li, F.; Han, L.; You, C.; and Wang, B. 2024. Segment anything in medical images. *Nature Communications*, 15(1): 654.

- Mandelli, S.; Cozzolino, D.; Cannas, E. D.; Cardenuto, J. P.; Moreira, D.; Bestagini, P.; Scheirer, W. J.; Rocha, A.; Verdoliva, L.; Tubaro, S.; et al. 2022. Forensic analysis of synthetically generated western blot images. *IEEE Access*, 10: 59919–59932.
- Moreira, D.; Cardenuto, J. P.; Shao, R.; Baireddy, S.; Cozzolino, D.; Gragnaniello, D.; Abd-Almageed, W.; Bestagini, P.; Tubaro, S.; Rocha, A.; et al. 2022. SILA: a system for scientific image analysis. *Scientific Reports*, 12(1): 18306.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis With Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10684–10695.
- Rublee, E.; Rabaud, V.; Konolige, K.; and Bradski, G. 2011. ORB: An efficient alternative to SIFT or SURF. In *2011 International conference on computer vision*, 2564–2571. Ieee.
- Ryu, S.-J.; Lee, M.-J.; and Lee, H.-K. 2010. Detection of copy-rotate-move forgery using Zernike moments. In *Information Hiding: 12th International Conference, IH 2010, Calgary, AB, Canada, June 28-30, 2010, Revised Selected Papers 12*, 51–65. Springer.
- Sabir, E.; Nandi, S.; Abd-Almageed, W.; and Natarajan, P. 2021. Biofors: A large biomedical image forensics dataset. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10963–10973.
- Sabir, E.; Nandi, S.; AbdAlmageed, W.; and Natarajan, P. 2022. MONet: Multi-Scale Overlap Network for Duplication Detection in Biomedical Images. In *2022 IEEE International Conference on Image Processing (ICIP)*, 3793–3797. IEEE.
- Shamshad, F.; Khan, S.; Zamir, S. W.; Khan, M. H.; Hayat, M.; Khan, F. S.; and Fu, H. 2023. Transformers in medical imaging: A survey. *Medical image analysis*, 88: 102802.
- Stern, A. M.; Casadevall, A.; Steen, R. G.; and Fang, F. C. 2014. Financial costs and personal consequences of research misconduct resulting in retracted publications. *Elife*, 3: e02956.
- Su, L.; Ma, X.; Zhu, X.; Niu, C.; Lei, Z.; and Zhou, J.-Z. 2025. Can we get rid of handcrafted feature extractors? sparsevit: Nonsemantics-centered, parameter-efficient image manipulation localization through spare-coding transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 7024–7032.
- Wang, J.; Chan, K. C.; and Loy, C. C. 2023. Exploring clip for assessing the look and feel of images. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 2555–2563.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.
- Wu, H.; Chen, S.; Chen, G.; Wang, W.; Lei, B.; and Wen, Z. 2022. FAT-Net: Feature adaptive transformers for automated skin lesion segmentation. *Medical image analysis*, 76: 102327.
- Wu, Y.; Abd-Almageed, W.; and Natarajan, P. 2018. Buster-net: Detecting copy-move image forgery with source/target localization. In *Proceedings of the European conference on computer vision (ECCV)*, 168–184.
- Wu, Y.; AbdAlmageed, W.; and Natarajan, P. 2019. Mantranet: Manipulation tracing network for detection and localization of image forgeries with anomalous features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9543–9552.
- Ying, Z.; Niu, H.; Gupta, P.; Mahajan, D.; Ghadiyaram, D.; and Bovik, A. 2020. From patches to pictures (PaQ-2-PiQ): Mapping the perceptual space of picture quality. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3575–3585.
- Yu, Q.; Zhang, J.; Zhang, H.; Wang, Y.; Lin, Z.; Xu, N.; Bai, Y.; and Yuille, A. 2021. Mask guided matting via progressive refinement network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1154–1163.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Zhang, S.; Xu, Y.; Usuyama, N.; Xu, H.; Bagga, J.; Tinn, R.; Preston, S.; Rao, R.; Wei, M.; Valluri, N.; Wong, C.; Tupini, A.; Wang, Y.; Mazzola, M.; Shukla, S.; Liden, L.; Gao, J.; Crabtree, A.; Piening, B.; Bifulco, C.; Lungren, M. P.; Naumann, T.; Wang, S.; and Poon, H. 2024. A Multimodal Biomedical Foundation Model Trained from Fifteen Million Image–Text Pairs. *NEJM AI*, 2(1).
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, 2223–2232.
- Zingarini, G.; Cozzolino, D.; Corvi, R.; Poggi, G.; and Verdoliva, L. 2024. M3DSYNTH: A dataset of medical 3D images with AI-generated local manipulations. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 13176–13180. IEEE.