

Weakly-Supervised Image Forgery Localization via Vision-Language Collaborative Reasoning Framework

Ziqi Sheng^{1,2}, Junyan Wu¹, Wei Lu^{1*}, Jiantao Zhou²

¹MoE Key Laboratory of Information Technology, School of Computer Science and Engineering, Sun Yat-sen University

²State Key Laboratory of Internet of Things for Smart City, Department of Computer and Information Science, University of Macau
shengzq3@mail.sysu.edu.cn, wujy298@mail2.sysu.edu.cn, luwei3@mail.sysu.edu.cn, jtzhou@um.edu.mo

Abstract

Image forgery localization aims to precisely identify tampered regions within images, but it commonly depends on costly pixel-level annotations. To alleviate this annotation burden, weakly supervised image forgery localization (WSIFL) has emerged, yet existing methods still achieve limited localization performance as they mainly exploit intra-image consistency clues and lack external semantic guidance to compensate for insufficient supervision information. In this paper, we propose ViLaCo, a vision-language collaborative reasoning framework that introduces auxiliary semantic supervision derived from pre-trained vision-language models (VLMs), enabling accurate pixel-level localization using only image-level labels. Specifically, we first employ a vision-language feature modeling network to jointly extract textual semantics and visual features by leveraging pre-trained VLMs. Next, an adaptive vision-language reasoning network aligns these features through mutual interactions, producing semantically aligned representations. Subsequently, these representations are passed into dual prediction heads, where the coarse head performs image-level classification and the fine head generates pixel-level localization masks, allowing the coarse-grained task to provide guidance for the fine-grained localization. Moreover, a contrastive patch consistency module is introduced to cluster tampered features while separating authentic ones, facilitating more reliable forgery discrimination. Extensive experiments on multiple public datasets demonstrate that ViLaCo substantially outperforms existing WSIFL methods, achieving state-of-the-art performance in both detection and localization accuracy.

Introduction

With the rapid development of AI-generated content technologies, image forgery has become increasingly common, and the era of “what you see is what you get” is gradually fading. The proliferation of malicious forged images is undermining public trust and posing serious threats to economic security and public safety. This drives the urgent need for reliable forgery localization methods that can accurately identify tampered regions. Although numerous approaches

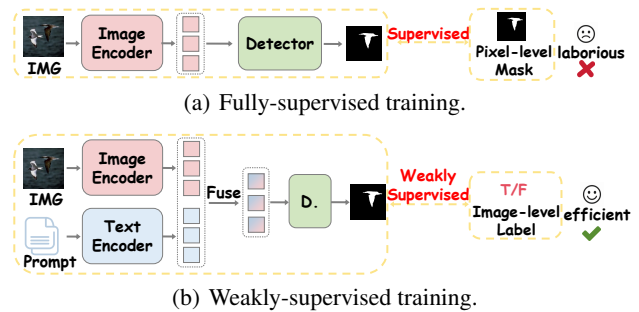


Figure 1: Difference between fully-supervised and weakly-supervised training strategies. (a) Fully-supervised methods use pixel-level masks for both training and prediction. (b) Weakly-supervised methods are trained only with binary image-level labels but are still required to predict pixel-level manipulation masks. Besides, the proposed method further leverages text prompts for extra supervision.

have been explored (Lou et al. 2025; Sheng et al. 2024), most methods rely on fully supervised learning, requiring costly pixel-level annotations during training (Fig. 1(a)). However, due to the rapid growth of forged images, collecting large-scale, high-quality pixel annotations is labor-intensive and costly, making it hard to apply in real-world settings.

To tackle these challenges, weakly supervised image forgery localization (WSIFL) has emerged as a promising direction. WSIFL methods aim to localize pixel-level manipulations using only image-level binary labels, typically leveraging contrastive learning and self-supervision to discover tamper-sensitive regions and enable fine-grained prediction. For instance, (Zhai et al. 2023) utilizes multi-scale consistency and inter-block correlations to discover the manipulation regions. (Li, Wen, and He 2025) employs block-level self-consistency and frame-level contrastive learning to distinguish consistent and inconsistent regions within images. Although these methods have achieved certain progress, they rely solely on internal image signals and lack external semantic guidance, which limits their ability to accurately identify forged regions.

*Corresponding Author

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

Meanwhile, vision-language models (Radford et al. 2021) have demonstrated remarkable capacity in aligning visual and textual modalities and improving downstream tasks like image classification and retrieval. However, in the field of WSIFL, its potential remains underexplored. In this paper, we aim to leverage the rich semantic alignment relationships between visual and textual modalities provided by pre-trained vision-language models (VLMs) to offer extra guidance for the WSIFL (Fig. 1(b)). Although VLMs provide a powerful external knowledge source, directly adapting them to the WSIFL task still faces three interrelated and progressive challenges. **(C1) Semantic-Artifact Mismatch.** VLMs are optimized for high-level semantic invariance, rendering them inherently insensitive to the subtle, low-level forgery signatures (e.g., splicing boundaries, noise inconsistencies) critical for manipulation detection. **(C2) Supervisory Granularity Gap.** The coarse, image-level supervision must be effectively translated into a precise, pixel-level guidance mechanism to localize specific artifact regions. **(C3) Prediction Ambiguity.** Weakly supervision permits trivial solutions, such as localizing only the most salient tampered pixels, which results in incomplete masks rather than a precise trace of the entire manipulated area.

To overcome these challenges, we propose ViLaCo, a vision-language collaborative reasoning framework for WSIFL. First, to address the semantic-artifact mismatch (C1), we introduce a vision-language feature modeling network. It re-purposes the VLM’s semantic representations by injecting low-level artifact awareness via a lightweight local-global spatial adapter, which is designed to model local inconsistencies and global dependencies. Second, to bridge the granularity gap (C2), we design an adaptive vision-language reasoning network featuring a novel forgery-aware aggregator. This module leverages the coarse, text-derived global forgery concepts to compute spatial guidance, which in turn selectively aggregates and enhances the visual features corresponding to the manipulated regions. Finally, to resolve prediction ambiguity (C3), we adopt a dual prediction strategy that decouples the task into two synergistic branches: a coarse branch for image-level classification and a fine branch for precise localization. Jointly optimizing these branches ensures the model identifies the forgery’s existence coarse while being driven to produce a complete and accurate pixel-level mask. Furthermore, a contrastive patch consistency loss is incorporated to further improve the discriminability of forgery features and enhance localization accuracy. Extensive experimental results over five public testing datasets demonstrate that ViLaCo significantly outperforms the state-of-the-art methods.

In summary, our contributions are as follows:

(1) We propose ViLaCo, a novel vision-language collaborative reasoning framework for WSIFL. By effectively leveraging semantic guidance from VLMs and a dual-branch architecture, ViLaCo simultaneously achieves effective image-level classification and accurate pixel-level localization using only image-level labels.

(2) We design a vision-language feature modeling network featuring a local-global spatial adapter (LGS adapter). This module successfully re-purposes the VLM’s high-level

semantic representations to be sensitive to the subtle, low-level artifacts indicative of forgery.

(3) We introduce an adaptive vision-language reasoning network featuring a novel forgery-aware aggregator. This module is responsible for translating coarse, text-derived global forgery concepts into precise, spatially-focused guidance by selectively aggregating and enhancing artifact-rich visual features.

(4) We employ a dual prediction strategy that decouples the task into synergistic coarse and fine branches. This structure, augmented by a contrastive patch consistency loss, enforces the generation of complete and spatially coherent forgery masks.

Related Works

Image Manipulation Localization

Fully supervised image manipulation localization methods exploit diverse forensic traces (Sheng et al. 2025; Triaridis and Mezaris 2024). Recent approaches also explore non-RGB domains, such as noise residuals (Li et al. 2024; Zeng et al. 2024; Guo et al. 2023) and frequency coefficients (Kwon et al. 2022; Wang et al. 2022), or utilize contrastive learning (Lou et al. 2025; Niloy, Bhaumik, and Woo 2023) to better capture subtle artifacts. Despite their effectiveness, these methods depend on laborious pixel-level annotations. To reduce this cost, weakly supervised (WSIFL) (Li, Wen, and He 2025; Zhai et al. 2023; Zhou et al. 2024) and auto-annotation (Qu et al. 2024) methods were developed. However, auto-annotation often requires multi-staged processing, while existing WSIFL methods are limited by their reliance on only internal image supervisory signals. In contrast, our method utilizes semantic supervision derived from pre-trained vision-language models, introducing semantic guidance beyond internal image clues to better uncover forgery traces.

Vision-Language Pre-training

Vision-language pre-training has shown strong ability to align visual and textual semantics by learning from large-scale image-text pairs. CLIP, as a representative model, demonstrates impressive generalization across classification, detection, captioning, and dense prediction tasks (Zhou et al. 2022a,b; Barraco et al. 2022; Rao et al. 2022). Recent studies further adapt these pre-trained models to specialized domains, including audio temporal forgery localization (Wu et al. 2025), video understanding (Luo et al. 2022), semantic segmentation (Kweon and Yoon 2024) and so on. Inspired by these advances, we explore leveraging CLIP’s rich vision-language representations for weakly supervised image forgery localization. Instead of relying purely on image-internal clues, our method brings in semantic supervision from pre-trained vision-language models, making weakly supervised forgery localization more precise and semantically informed.

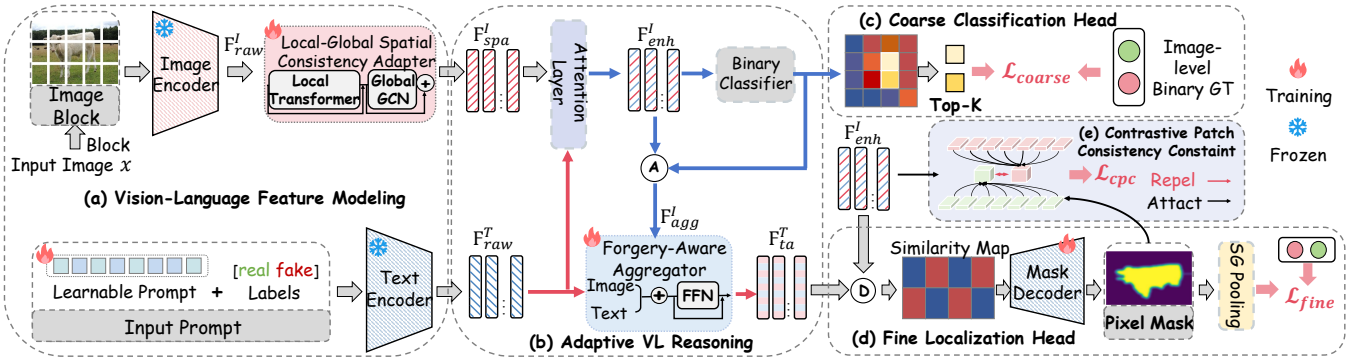


Figure 2: The proposed ViLaCo framework consists of (a) vision-language feature modeling, (b) adaptive vision-language reasoning, (c) coarse classification head, (d) fine localization head, and (e) a contrastive patch consistency constraint.

Method

Problem Definition

The weakly-supervised image forgery localization (WSIFL) task supposes that only image-level labels are available during the training stage, and encourages the models to predict whether each pixel is tampered at the inference stage. Mathematically, given a set of samples $\mathcal{X}, \mathcal{Y}, \mathcal{M}$, where $\mathcal{X}, \mathcal{Y}, \mathcal{M}$ denote the sets of image, image-level binary label, and pixel-level localization mask, respectively. For each input image $x \in \mathcal{R}^{H \times W \times C}$, it has two corresponding labels, namely, y and m , with m only for the inference stage. Here $y \in \{0, 1\}$ and $y = 1$ indicates that x is a tampered image, and $m \in \{0, 1\}^{H \times W \times 1}$, $m = 1$ indicated the tampered regions.

Overview

As illustrated in Fig. 2, we propose ViLaCo, the vision-language collaborative reasoning framework for weakly supervised image forgery localization (WSIFL). Unlike previous weakly supervised methods that relied solely on internal image signals, ViLaCo re-purposes the representations of a pre-trained vision-language model, transforming its high-level semantic knowledge into precise, artifact-aware guidance for pixel-level localization under weak supervision. The framework consists of four key components:

(i) Vision-Language Feature modeling. The input image is partitioned into blocks and passed through a frozen image encoder to extract raw visual features \mathcal{F}_{raw}^I , which are subsequently refined by a learnable local-global spatial adapter (LGS-adapter) to capture spatial forgery clues, yielding \mathcal{F}_{spa}^I . For textual input, a learnable prompt combined with binary class labels is processed by a frozen text encoder to generate text features \mathcal{F}_{raw}^T .

(ii) Adaptive Vision-Language Reasoning. To fully leverage cross-modal information, ViLaCo adaptively fuses \mathcal{F}_{raw}^T and \mathcal{F}_{spa}^I , producing text-enhanced visual features \mathcal{F}_{enh}^I via an attention layer. Meanwhile, a forgery-aware aggregator jointly models visual-text interactions to obtain a tampering-aware textual embedding \mathcal{F}_{ta}^T . This adaptive reasoning facilitates discriminative feature learning for subsequent predictions.

(iii) Dual-Branch Coarse-to-Fine Architecture. To bridge the gap between image-level supervision and pixel-level localization, ViLaCo introduces a dual prediction head design. The coarse classification head estimates the forgery likelihood by aggregating top- K suspicious patches, enabling robust binary classification. Subsequently, the fine localization head constructs a similarity map between \mathcal{F}_{enh}^I and \mathcal{F}_{ta}^T , which is decoded into a pixel-level mask via a mask decoder. Besides, a soft-gated pooling layer further converts the mask into an auxiliary binary prediction, thereby enabling supervision using binary labels.

(iv) Contrastive Patch Consistency Constraint. To enhance the discriminability of forgery features without pixel-level ground truth, a novel contrastive consistency constraint \mathcal{L}_{cpc} is proposed. It is designed to ensure that patches with similar forged clues in \mathcal{F}_{enh}^I are pulled closer together, while patches with significant differences are pushed apart. This mechanism encourages features to cluster around manipulated regions, thereby improving mask quality.

Vision-Language Feature Modeling

The vision-language feature modeling network aims to provide reliable and modality-specific representations of both image and text inputs, serving as the foundation for subsequent cross-modal reasoning and localization. For the visual modality, a frozen image encoder is employed to obtain raw image embeddings \mathcal{F}_{raw}^I . These raw embeddings are then refined by the local-global spatial adapter (LGS-adapter), which re-purposes the features to shift their focus from high-level semantic invariance to low-level forgery traces. Meanwhile, for the textual modality, we directly concatenate trainable prompt embeddings with the fixed label tokens (“real” and “fake”), forming adaptive text representations.

Local-Global Spatial Consistency Adapter To effectively encode spatial forgery clues, we propose an LGS-adapter to model both local inconsistencies and global structural dependencies within the encoded image feature. LGS-adapter sequentially integrates a local transformer encoder and a lightweight graph convolutional network (GCN) (Chen et al. 2020), enabling spatially locally and globally aware feature learning.

Given the patch-level visual features $\mathcal{F}_{raw}^I \in \mathbb{R}^{n \times d}$ extracted from a frozen image encoder, where n denotes the number of image patches and d is the feature dimension, we first capture local spatial interactions X_l using a windowed transformer encoder. Unlike standard transformers that compute self-attention globally, this layer restricts attention computation to non-overlapping and partially overlapping spatial windows, focusing on nearby patch interactions without cross-window message passing. This local design simulates convolutional receptive fields, enhancing the network’s sensitivity to subtle manipulations while significantly reducing computational cost.

While the local transformer is effective for patch-level forgery clues, it lacks the ability to model long-range spatial dependencies that may arise in complex tampering scenarios. To complement this, we incorporate a global reasoning step via a lightweight GCN module. We construct two adjacency matrices, H_{sim} for capturing pairwise patch similarities and H_{dis} for modeling relative spatial distances between patches. Both matrices are row-normalized using softmax to ensure stable feature aggregation. The GCN propagates information across all patches to produce globally consistent features:

$$X_g = \text{GELU}([\text{Softmax}(H_{sim}); \text{Softmax}(H_{dis})] X_l W), \quad (1)$$

where W is a learnable transformation matrix. This global aggregation enriches each patch feature with context from distant regions, allowing the model to capture spatially coherent manipulation patterns even when tampered areas are spatially separated. By adaptively combining local attention and global graph reasoning, the LGS-adapter produces forgery-aware features \mathcal{F}_{spa}^I that are both locally discriminative and globally consistent, providing a strong foundation for subsequent weakly supervised localization.

Learnable Prompt In weakly supervised image forgery localization, the textual labels (“real” and “fake”) are insufficient to fully express the complex semantic information associated with image manipulations. Such minimal labels can limit the transferability of textual embeddings and hinder effective vision-language interaction. Inspired by recent advances in prompt learning (Zhou et al. 2022a), we enhance the textual representation by introducing a set of learnable context tokens, forming a more adaptable and forgery-aware prompt.

Specifically, the discrete binary label is first tokenized using a pre-trained CLIP tokenizer to obtain an initial class token $t_{init} = \text{Tokenizer}(\text{label})$, where $\text{label} \in \{\text{real}, \text{fake}\}$. To construct a forgery-aware prompt, we concatenate t_{init} with a sequence of learnable context tokens $\{c_1, \dots, c_l\}$, yielding:

$$t_p = \{c_1, \dots, t_{init}, \dots, c_l\}. \quad (2)$$

The class token is placed at the center of the sequence to promote balanced contextualization from both directions. The combined prompt is then fused with positional embeddings to preserve token order and spatial relevance. Finally, the text encoder of CLIP processes t_p and produces the enhanced textual embedding $t_{out} \in \mathbb{R}^d$. By adaptively learning

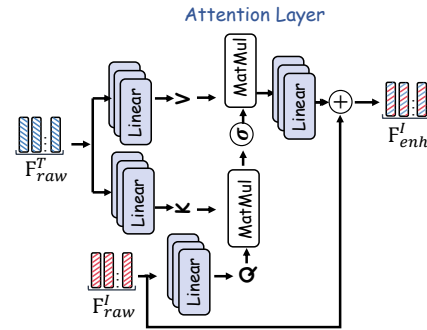


Figure 3: Illustration of the structure attention layer, where textual features F_{raw}^T guide the enhancement of visual features F_{raw}^I .

the forgery-aware prompt t_p , this module enables the textual representation to better align with diverse manipulation patterns, thus facilitating effective cross-modal reasoning and improving the performance of WSIFL.

Adaptive Vision-Language Reasoning

We design the adaptive vision-language reasoning network to translate coarse, text-derived forgery concepts into precise, spatially-focused guidance. This is achieved through two main functions: (i) adaptively enhancing artifact-sensitive visual features using clues from the textual embeddings, and (ii) integrating the resulting cross-modal information to construct a similarity map for fine-grained forgery localization.

As illustrated in Fig. 3, given the raw textual embedding F_{raw}^T and the patch-level visual features F_{raw}^I obtained from the LGS-adapter, we employ a text-guided attention layer to selectively enhance visual representations. Concretely, F_{raw}^I serves as a query to attend to F_{raw}^T , allowing the model to amplify patches that are semantically related to tampering clues while suppressing irrelevant regions to obtain the text-enhanced visual features F_{enh}^I . This step enriches visual representations with manipulation-aware semantics, thereby enhancing the discriminative nature of feature F_{enh}^I .

To further strengthen cross-modal alignment, we introduce a forgery-aware aggregator that refines textual embeddings based on aggregated visual context. Specifically, we first summarize the enhanced visual features F_{enh}^I into a compact visual context vector F_{agg}^I via a soft attention pooling mechanism that highlights the most manipulation-indicative patches. This aggregated feature F_{agg}^I is then fused with the raw textual embedding through a feed-forward network (FFN):

$$F_{ta}^T = \text{FFN}(F_{agg}^I + F_{raw}^T) + F_{agg}^I + F_{raw}^T, \quad (3)$$

where F_{ta}^T denotes the tampering-aware textual embedding. The forgery-aware aggregator enables bidirectional interaction between modalities, allowing textual features to dynamically adapt to image-specific manipulation patterns. This tampering-aware textual embedding F_{ta}^T , together with F_{enh}^I , provides strong discriminative signals for downstream coarse classification and fine-grained localization.

Dual-Branch Coarse-to-Fine Architecture

To bridge the gap between image-level supervision and pixel-level tampering localization, we design a dual-branch coarse-to-fine architecture. This component decomposes the weakly supervised localization task into two complementary heads: a coarse classification head for effective image-level forgery detection and a fine localization head for accurate pixel-level mask prediction. The two heads work in a collaborative manner, enabling the model to gradually improve its ability to detect tampered regions.

Coarse Classification Head Given the text-enhanced visual features F_{enh}^I obtained from the adaptive vision-language reasoning module, we first predict the tampering probability of each image patch through a binary classifier. To focus on the most suspicious regions, we select the top- K patches with the highest tampering probability and aggregate their scores to form an image-level prediction \hat{y}_{coarse} . The top- K pooling strategy simulates the weakly supervised setting, where only a few localized patches are manipulated while the remaining are authentic. The coarse classification loss is defined as a binary cross-entropy (BCE) loss:

$$\mathcal{L}_{coarse} = -[y \log(\hat{y}_{coarse}) + (1-y) \log(1-\hat{y}_{coarse})], \quad (4)$$

where $y \in \{0, 1\}$ is the ground-truth image-level label. This branch ensures that the network can effectively distinguish forged images from pristine ones.

Fine Localization Head To achieve pixel-level localization, we design a fine localization head that decodes patch-level features into a pixel-level manipulation mask. Specifically, the enhanced visual features F_{enh}^I and the tamper-resistant text embeddings F_{ta}^T are combined via a dot product operation to obtain a cross-modal similarity map, which is then input into a masking decoder to predict pixel-level masks \hat{M} .

To leverage weak image-level supervision while learning spatially discriminative features, we introduce an adaptive soft-gated pooling (SG pooling) layer on top of \hat{M} . Unlike conventional max or average pooling, which either focus excessively on peak responses or dilute discriminative pixel clues, SG pooling employs a differentiable gating mechanism with learnable threshold and temperature parameters. The mechanism adaptively assigns higher weights to manipulated pixels and suppresses background responses, enabling robust aggregation of pixel-level predictions into an image-level score \hat{y}_{fine} . By maintaining differentiability, SG pooling facilitates end-to-end optimization and allows the network to emphasize manipulation-relevant pixels automatically. The localization loss for this branch is defined as a binary cross-entropy loss:

$$\mathcal{L}_{fine} = -[y \log(\hat{y}_{fine}) + (1-y) \log(1-\hat{y}_{fine})], \quad (5)$$

where $y \in \{0, 1\}$ is the image-level ground-truth label.

By jointly optimizing \mathcal{L}_{coarse} and \mathcal{L}_{fine} , the dual-branch architecture enables the model to first coarsely determine whether an image contains manipulations and then further identify pixel-level tampered regions.

Contrastive Patch Consistency Constraint

Although the dual-branch architecture provides predictions from coarse to fine, the absence of GT pixel-level labels still hinders the distinction between real and fake parts of image features. Therefore, based on the idea of self-supervised learning, we propose a contrastive patch consistency (CPC) constraint, which encourages parts with similar forgery clues to cluster together while pushing away parts from the true region. The constraint enhances the discriminative ability of the learned image representations and improves the quality of the predicted forgery mask.

Based on the patch-level enhanced visual feature F_{enh}^I , we assign a label to each patch based on the predicted manipulation mask \hat{M} . Specifically, patches with responses above a threshold τ_{fg} are treated as *tampered*, while those below τ_{bg} are treated as *authentic*. This labeling allows us to build positive and negative patch pairs without requiring ground-truth masks.

Concretely, for a tampered patch feature f_i^{tam} and an authentic patch feature f_j^{real} , we compute the similarity scores via normalized dot products. The \mathcal{L}_{cpc} is formulated as:

$$\mathcal{L}_{cpc} = -\frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} \left[\log \frac{\exp(\text{sim}(f_i^{tam}, f_j^{tam})/\gamma)}{\sum_k \exp(\text{sim}(f_i^{tam}, f_k)/\gamma)} + \log \frac{\exp(\text{sim}(f_j^{real}, f_i^{real})/\gamma)}{\sum_k \exp(\text{sim}(f_j^{real}, f_k)/\gamma)} \right], \quad (6)$$

where \mathcal{P} denotes the set of positive patch pairs, γ is a temperature parameter, and $\text{sim}(\cdot, \cdot)$ represents cosine similarity. The first term pulls together features from tampered patches, while the second term aligns authentic patches. Negatives are sampled from patches of the opposite type, thereby pushing apart dissimilar features.

Objective Function

The entire framework is trained end-to-end under weak supervision by jointly optimizing the losses from all predicted heads. The total loss is defined as:

$$\mathcal{L} = \mathcal{L}_{coarse} + \mathcal{L}_{fine} + \lambda_{ccs}(t) \mathcal{L}_{cpc}, \quad (7)$$

where \mathcal{L}_{coarse} and \mathcal{L}_{fine} are the binary classification losses from the coarse classification head and fine localization head, respectively. \mathcal{L}_{cpc} is the contrastive patch consistency constraint that regularizes patch-level feature learning. Unlike a fixed weighting, we adopt a warm-up scheduling strategy (Goyal et al. 2017) for the consistency term:

$$\lambda_{ccs}(t) = \begin{cases} 0, & t < T_w, \\ 1 - \exp\left(-\frac{t-T_w}{T_{total}-T_w}\right), & t \geq T_w, \end{cases} \quad (8)$$

where t is the current training epoch, T_w is the warm-up starting epoch, and T_{total} is the total number of training epochs. This schedule gradually activates the contrastive patch consistency constraint after an initial warm-up period, allowing the network to first learn stable image-level discrimination before enforcing spatial consistency. All modules of ViLaCo are optimized jointly via stochastic gradient descent with backpropagation.

	Baselines	Pixel-Level F1						Combined F1				
		CASIAv1	Columbia	Coverage	IMD2020	NIST16	AVG	CASIAv1	Columbia	Coverage	IMD2020	AVG
Un.	NOI1	0.157	0.311	0.205	0.124	0.089	0.190	0.000	0.000	0.000	0.000	0.000
	CFA1	0.140	0.320	0.188	0.111	0.106	0.188	0.000	0.000	0.000	0.000	0.000
Fully-supervised	H-LSTM	0.154	0.130	0.163	0.195	0.354	0.176	0.000	0.004	0.000	0.000	0.001
	ManTra-Net	0.155	0.364	0.286	0.122	0.000	0.185	0.000	0.000	0.000	0.000	0.000
	RRU-Net	0.225	0.452	0.189	0.232	0.265	0.273	0.023	0.000	0.000	0.000	0.006
	CR-CNN	0.405	0.436	0.291	-	0.238	-	0.382	0.413	0.181	-	-
	GSR-Net	0.387	0.613	0.285	0.175	0.283	0.349	0.042	0.042	0.000	0.026	0.028
	SPAN	0.184	0.487	0.172	0.170	0.221	0.214	0.000	0.000	0.000	0.000	0.000
	CAT-Net	0.276	0.352	0.134	0.102	0.138	0.200	0.345	0.406	0.149	0.144	0.261
	MVSS-Net	0.452	0.638	0.453	0.260	0.292	0.419	0.566	0.711	0.317	0.300	0.474
	IF-OSN	0.686	0.728	0.743	0.576	0.645	0.676	0.857	0.904	0.678	0.547	0.747
Weakly-supervised	MIL-FCN	0.117	0.089	0.121	0.097	0.024	0.090	0.193	0.141	0.118	0.131	0.146
	MIL-FCN+WACL	0.172	0.270	0.178	0.193	0.110	0.185	0.280	0.386	0.268	0.252	0.296
	Araslanov	0.112	0.102	0.127	0.094	0.026	0.092	0.194	0.140	0.133	0.046	0.125
	Araslanov+WACL	0.153	0.362	0.201	0.173	0.099	0.198	0.250	0.414	0.255	0.159	0.270
	EdgeCAM	0.338	0.470	0.262	0.242	0.254	0.313	0.476	0.573	0.297	0.347	0.423
	WSCCL	<u>0.347</u>	<u>0.273</u>	<u>0.301</u>	<u>0.265</u>	<u>0.159</u>	<u>0.269</u>	0.475	0.302	<u>0.427</u>	<u>0.365</u>	0.388
	MRL-Net	<u>0.347</u>	<u>0.534</u>	0.213	0.248	0.113	0.265	<u>0.495</u>	<u>0.603</u>	0.316	0.348	<u>0.441</u>
	ViLaCo	0.491	0.536	0.319	0.365	0.267	0.373	0.632	0.714	0.568	0.456	0.593

Table 1: Comparison with unsupervised (Un.), fully-supervised and weakly-supervised methods on pixel-level manipulation localization pF1 score and the combined F1 score between I-F1 and P-F1. The best and the second best results in weakly-supervised methods are noted with **bolded** and underlined respectively.

Experiments

Setup

Datasets For consistency and fairness, we follow the settings of previous weakly supervised image localization methods (Zhai et al. 2023; Li, Wen, and He 2025). Our experiments are trained on the CASIAv2 dataset (Dong, Wang, and Tan 2013), with in-dataset testing on CASIAv1 and cross-dataset testing on Columbia (Hsu and Chang 2006), COVER (Wen et al. 2016), NIST16 (Guan et al. 2019), and IMD20 (Novozamsky, Mahdian, and Saic 2020).

Evaluation Metrics We assess localization performance using pixel-level F1 (P-F1) for manipulated regions and combined F1 (C-F1) for overall accuracy, both computed with a fixed threshold of 0.5. Image-level detection is evaluated using image F1 (I-F1).

Implementation Details Our network adopts frozen image and text encoders from the pre-trained CLIP (ViT-B/16), with transformer-based FFN layers where ReLU activations are replaced by GELU. All input images are resized to 256×256 and augmented via standard cropping and flipping. The patch size is set to 8×8 . The model is implemented in PyTorch and trained on a single NVIDIA RTX 4090 GPU using the AdamW optimizer (Loshchilov and Hutter 2017) with a batch size of 32, an initial learning rate of 0.0001, and a total of 100 epochs.

Comparison with State-of-the-Art

In this section, we compare ViLaCo’s image-level detection and pixel-level localization performance with 18 existing methods. Unsupervised: CFA1 (Ferrara et al. 2012), NOI1 (Mahdian and Saic 2009); fully supervised: H-LSTM (Bappy et al. 2019), ManTra-Net (Wu, AbdAlmageed, and Natarajan 2019), RRU-Net (Bi et al. 2019), CR-CNN (Yang

et al. 2019), GSR-Net (Zhou et al. 2020), SPAN (Hu et al. 2020), CAT-Net (Kwon et al. 2022), FCN+DA (Chen et al. 2021), MVSS-Net (Dong et al. 2022), IF-OSN (Wu et al. 2022); weakly supervised: MIL-FCN (Pathak et al. 2014), Araslanov (Araslanov and Roth 2020), WACL (Zhai et al. 2023), EdgeCAM (Zhou et al. 2024), WSCCL (Bai 2025), MRL-Net (Li, Wen, and He 2025).

Pixel-Level Localization Comparisons The pixel-level localization results are presented in Tab. 1, ViLaCo is compared with unsupervised, fully supervised, and weakly supervised baselines across five datasets. ViLaCo achieves state-of-the-art results among weakly-supervised approaches, significantly outperforming counterparts in both pixel-level F1 and combined F1 scores. Moreover, our approach shows strong competitiveness compared to fully-supervised methods; notably, ViLaCo improves over MVSS-Net on average Combined F1 by 11.9%. These observations demonstrate that ViLaCo not only achieves accurate pixel-level localization under limited supervision but also exhibits superior generalization and robustness to diverse manipulation scenarios.

Image-Level Detection Comparisons To assess image-level detection performance, we compare ViLaCo with state-of-the-art baselines (Tab. 2). ViLaCo achieves the highest average I-F1 score of 0.776, consistently outperforming weakly supervised methods and remaining competitive with fully supervised approaches, further validating its effectiveness under limited supervision.

Qualitative Results

We further provide qualitative results to visually illustrate the effectiveness of ViLaCo in localizing manipulated regions. As depicted in Fig. 4, our method demonstrates superior performance compared to both unsupervised and ex-

	Baselines	CASIAv1	Columbia	COVER	IMD20	AVG
Un.	NOII	0.000	0.000	0.000	0.000	0.000
	CFAI	0.000	0.000	0.000	0.000	0.000
Fully.	H-LSTM	0.000	0.002	0.000	0.000	0.001
	ManTra-Net	0.000	0.000	0.000	0.000	0.000
	RRU-Net	0.001	0.000	0.000	0.000	0.000
	CR-CNN	0.361	0.392	0.131	0.200	0.271
	GSR-Net	0.022	0.022	0.000	0.014	0.019
	SPAN	0.000	0.000	0.000	0.000	0.000
	CAT-Net	0.459	0.505	0.169	0.229	0.157
	FCN+DA	0.775	0.481	0.180	0.182	0.404
MVSS-Net	0.758	0.802	0.244	0.355	0.534	
Weakly.	MIL-FCN	0.553	0.338	0.115	0.205	0.303
	MIL-FCN+WSDL	0.738	0.680	0.544	0.360	0.580
	Araslanov	0.496	0.140	0.140	0.219	0.270
	Araslanov+WSDL	0.679	0.483	0.348	0.316	0.456
	EdgeCAM	0.806	0.733	0.343	<u>0.613</u>	0.624
	MRL-Net	0.866	0.975	0.610	0.585	0.762
	ViLaCo	0.880	<u>0.917</u>	0.630	0.675	0.776

Table 2: Comparison of state-of-the-art methods for image-level manipulation detection across multiple datasets, evaluated by image-level F1 score. The first and second rankings are shown in **bolded** and underlined respectively.

ID	LOSS			CASIA		IMD20	
	\mathcal{L}_{coarse}	\mathcal{L}_{fine}	\mathcal{L}_{cpc}	PF1	IF1	PF1	IF1
1	-	✓	✓	0.842	0.410	0.620	0.305
2	✓	-	✓	0.790	0.365	0.575	0.270
3	✓	✓	-	0.860	0.445	0.640	0.320
4	✓	✓	✓	0.880	0.491	0.675	0.365

Table 3: Ablation study of the proposed \mathcal{L}_{coarse} , \mathcal{L}_{fine} and \mathcal{L}_{cpc} in terms of F1 score. The bold mark best performance

isting weakly-supervised methods. Compared with existing weakly-supervised methods, ViLaCo achieves significantly better localization performance and reaches a quality comparable to fully-supervised methods. This result highlights ViLaCo’s capability to produce clearer and more accurate pixel-level localization masks by leveraging vision-language alignment.

Fig. 5 shows the impact of block size and prediction heads. An 8×8 block size yields the best accuracy-efficiency trade-off. Furthermore, the convergence of coarse and fine heads confirms their collaborative contribution to stable detection.

Ablation Study

To investigate the contribution of each loss function in ViLaCo, we perform an ablation study by selectively removing \mathcal{L}_{coarse} , \mathcal{L}_{fine} , and \mathcal{L}_{cpc} , and report the results in Tab. 3.

When \mathcal{L}_{coarse} is removed (ID 1), the model relies only on fine-branch supervision, leading to a noticeable drop in detection accuracy (IF1 decreases from 0.491 to 0.410 on CASIA). Removing \mathcal{L}_{fine} (ID 2) causes similar degradation, indicating both branches play complementary roles in localization. Excluding \mathcal{L}_{cpc} (ID 3) also reduces performance, notably on IMD20 where IF1 drops from 0.365 to 0.320, showing \mathcal{L}_{cpc} effectively enhances forgery distinction. Finally, combining all losses (ID 4) yields the best results,

demonstrating that joint optimization is crucial for accurate pixel- and image-level localization.




Figure 4: Qualitative comparison of ViLaCo on DEFACTO datasets with (a) unsupervised, (b) weakly-supervised and (c) fully-supervised methods.

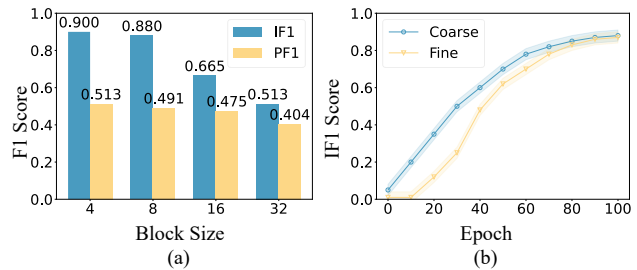


Figure 5: Effect of (a) block size and (b) the prediction head on detection and localization performance.

Conclusion

In this work, we presented ViLaCo, the vision-language collaborative reasoning framework for weakly supervised image forgery localization. Unlike prior WSIFL approaches that solely rely on intra-image consistency clues, ViLaCo introduces auxiliary semantic supervision derived from pre-trained vision-language models, enabling precise localization using only image-level annotations. The proposed architecture progressively integrates this semantic knowledge through three key designs: a vision-language feature modeling network with a local-global spatial adapter to capture forgery-specific clues, an adaptive vision-language reasoning network that fuses textual semantics and visual features to highlight manipulated regions, and a coarse-to-fine dual prediction mechanism enhanced by a contrastive patch consistency loss to refine mask quality. Extensive experiments across multiple benchmarks demonstrate that ViLaCo significantly outperforms existing weakly supervised methods and achieves competitive performance compared to fully supervised counterparts.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (No.62441237).

References

- Araslanov, N.; and Roth, S. 2020. Single-stage semantic segmentation from image labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4253–4262.
- Bai, R. 2025. Weakly-supervised cross-contrastive learning network for image manipulation detection and localization. *Knowledge-Based Systems*, 310: 113033.
- Bappy, J. H.; Simons, C.; Nataraj, L.; Manjunath, B.; and Roy-Chowdhury, A. K. 2019. Hybrid LSTM and encoder–decoder architecture for detection of image forgeries. *IEEE Transactions on Image Processing*, 28(7): 3286–3300.
- Barraco, M.; Cornia, M.; Cascianelli, S.; Baraldi, L.; and Cucchiara, R. 2022. The unreasonable effectiveness of CLIP features for image captioning: an experimental analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4662–4670.
- Bi, X.; Wei, Y.; Xiao, B.; and Li, W. 2019. RRU-Net: The ringed residual U-Net for image splicing forgery detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition workshops*, 0–0.
- Chen, M.; Wei, Z.; Huang, Z.; Ding, B.; and Li, Y. 2020. Simple and deep graph convolutional networks. In *International Conference on Machine Learning*, 1725–1735.
- Chen, X.; Dong, C.; Ji, J.; Cao, J.; and Li, X. 2021. Image manipulation detection by multi-view multi-scale supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14185–14193.
- Dong, C.; Chen, X.; Hu, R.; Cao, J.; and Li, X. 2022. Mvssnet: Multi-view multi-scale supervised networks for image manipulation detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3): 3539–3553.
- Dong, J.; Wang, W.; and Tan, T. 2013. Casia image tampering detection evaluation database. In *IEEE China Summit and International Conference on Signal and Information Processing*, 422–426.
- Ferrara, P.; Bianchi, T.; De Rosa, A.; and Piva, A. 2012. Image forgery localization via fine-grained analysis of CFA artifacts. *IEEE Transactions on Information Forensics and Security*, 7(5): 1566–1577.
- Goyal, P.; Dollár, P.; Girshick, R.; Noordhuis, P.; Wesolowski, L.; Kyrola, A.; Tulloch, A.; Jia, Y.; and He, K. 2017. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*.
- Guan, H.; Kozak, M.; Robertson, E.; Lee, Y.; Yates, A. N.; Delgado, A.; Zhou, D.; Kheyrkhan, T.; Smith, J.; and Fiscus, J. 2019. MFC datasets: Large-scale benchmark datasets for media forensic challenge evaluation. In *IEEE Winter Applications of Computer Vision Workshops*, 63–72.
- Guo, X.; Liu, X.; Ren, Z.; Grosz, S.; Masi, I.; and Liu, X. 2023. Hierarchical fine-grained image forgery detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3155–3165.
- Hsu, Y.-F.; and Chang, S.-F. 2006. Detecting image splicing using geometry invariants and camera characteristics consistency. In *IEEE International Conference on Multimedia and Expo*, 549–552.
- Hu, X.; Zhang, Z.; Jiang, Z.; Chaudhuri, S.; Yang, Z.; and Nevatia, R. 2020. SPAN: Spatial pyramid attention network for image manipulation localization. In *European Conference on Computer Vision*, 312–328.
- Kweon, H.; and Yoon, K.-J. 2024. From sam to cams: Exploring segment anything model for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19499–19509.
- Kwon, M.-J.; Nam, S.-H.; Yu, I.-J.; Lee, H.-K.; and Kim, C. 2022. Learning jpeg compression artifacts for image manipulation detection and localization. *International Journal of Computer Vision*, 130(8): 1875–1895.
- Li, J.; Wen, Y.; and He, L. 2025. M²RL-Net: Multi-View and Multi-Level Relation Learning Network for Weakly-Supervised Image Forgery Detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 4743–4751.
- Li, S.; Ma, W.; Guo, J.; Xu, S.; Li, B.; and Zhang, X. 2024. Unionformer: Unified-learning transformer with multi-view representation for image manipulation detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12523–12533.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Lou, Z.; Cao, G.; Guo, K.; Yu, L.; and Weng, S. 2025. Exploring multi-view pixel contrast for general and robust image forgery localization. *IEEE Transactions on Information Forensics and Security*.
- Luo, H.; Ji, L.; Zhong, M.; Chen, Y.; Lei, W.; Duan, N.; and Li, T. 2022. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508: 293–304.
- Mahdian, B.; and Saic, S. 2009. Using noise inconsistencies for blind image forensics. *Image and vision computing*, 27(10): 1497–1503.
- Niloy, F. F.; Bhaumik, K. K.; and Woo, S. S. 2023. CFL-Net: Image forgery localization using contrastive learning. In *Proceedings of the IEEE/CVF winter Conference on Applications of Computer Vision*, 4642–4651.
- Novozamsky, A.; Mahdian, B.; and Saic, S. 2020. IMD2020: A large-scale annotated dataset tailored for detecting manipulated images. In *Proceedings of the IEEE/CVF winter Conference on Applications of Computer Vision Workshops*, 71–80.
- Pathak, D.; Shelhamer, E.; Long, J.; and Darrell, T. 2014. Fully convolutional multi-class multiple instance learning. *arXiv preprint arXiv:1412.7144*.

- Qu, C.; Zhong, Y.; Liu, C.; Xu, G.; Peng, D.; Guo, F.; and Jin, L. 2024. Towards modern image manipulation localization: A large-scale dataset and novel methods. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10781–10790.
- 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.
- Rao, Y.; Zhao, W.; Chen, G.; Tang, Y.; Zhu, Z.; Huang, G.; Zhou, J.; and Lu, J. 2022. Denseclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18082–18091.
- Sheng, Z.; Lu, W.; Luo, X.; Zhou, J.; and Cao, X. 2025. SUMI-IFL: An Information-Theoretic Framework for Image Forgery Localization with Sufficiency and Minimality Constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 720–728.
- Sheng, Z.; Qu, Z.; Lu, W.; Cao, X.; and Huang, J. 2024. DiRLoc: Disentanglement Representation Learning for Robust Image Forgery Localization. *IEEE Transactions on Dependable and Secure Computing*.
- Triaridis, K.; and Mezaris, V. 2024. Exploring multi-modal fusion for image manipulation detection and localization. In *International Conference on Multimedia Modeling*, 198–211.
- Wang, J.; Wu, Z.; Chen, J.; Han, X.; Shrivastava, A.; Lim, S.-N.; and Jiang, Y.-G. 2022. Objectformer for image manipulation detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2364–2373.
- Wen, B.; Zhu, Y.; Subramanian, R.; Ng, T.-T.; Shen, X.; and Winkler, S. 2016. COVERAGE—A novel database for copy-move forgery detection. In *IEEE International Conference on Image Processing*, 161–165.
- Wu, H.; Zhou, J.; Tian, J.; Liu, J.; and Qiao, Y. 2022. Robust image forgery detection against transmission over online social networks. *IEEE Transactions on Information Forensics and Security*, 17: 443–456.
- Wu, J.; Xu, W.; Lu, W.; Luo, X.; Yang, R.; and Guo, S. 2025. Weakly-supervised Audio Temporal Forgery Localization via Progressive Audio-language Co-learning Network. *arXiv preprint arXiv:2505.01880*.
- 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.
- Yang, C.; Li, H.; Lin, F.; Jiang, B.; and Zhao, H. 2019. Constrained R-CNN: A general image manipulation detection model. *arXiv preprint arXiv:1911.08217*.
- Zeng, K.; Cheng, R.; Tan, W.; and Yan, B. 2024. Mgq-former: Mask-guided query-based transformer for image manipulation localization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 6944–6952.
- Zhai, Y.; Luan, T.; Doermann, D.; and Yuan, J. 2023. Towards generic image manipulation detection with weakly-supervised self-consistency learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 22390–22400.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022a. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9): 2337–2348.
- Zhou, P.; Chen, B.-C.; Han, X.; Najibi, M.; Shrivastava, A.; Lim, S.-N.; and Davis, L. 2020. Generate, segment, and refine: Towards generic manipulation segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 13058–13065.
- Zhou, X.; Girdhar, R.; Joulin, A.; Krähenbühl, P.; and Misra, I. 2022b. Detecting twenty-thousand classes using image-level supervision. In *European Conference on Computer Vision*, 350–368.
- Zhou, Y.; Wang, H.; Zeng, Q.; Zhang, R.; and Meng, S. 2024. Exploring weakly-supervised image manipulation localization with tampering edge-based class activation map. *Expert Systems with Applications*, 249: 123501.