

Mesoscopic Insights: Orchestrating Multi-Scale & Hybrid Architecture for Image Manipulation Localization

Xuekang Zhu^{1,2*}, Xiaochen Ma^{1,2*}, Lei Su^{1,2}, Zhuohang Jiang³, Bo Du^{1,2}, Xiwen Wang^{1,2}, Zeyu Lei^{1,2}, Wentao Feng^{1,2}, Chi-Man Pun⁴, Ji-Zhe Zhou^{1,2†}

¹College of Computer Science, Sichuan University

²Engineering Research Center of Machine Learning and Industry Intelligence, Ministry of Education of China

³The Hong Kong Polytechnic University

⁴Computer and Information Science, Faculty of Science and Technology, University of Macau

Abstract

The mesoscopic level serves as a bridge between the macroscopic and microscopic worlds, addressing gaps overlooked by both. Image manipulation localization (IML), a crucial technique to pursue truth from fake images, has long relied on low-level (microscopic-level) traces. However, in practice, most tampering aims to deceive the audience by altering image semantics. As a result, manipulation commonly occurs at the object level (mesoscopic level), which is equally important as microscopic traces. Therefore, integrating these two levels into the mesoscopic level presents a new perspective for IML research. Inspired by this, our paper explores how to simultaneously construct mesoscopic representations of micro and macro information for IML and introduces the Mesorch architecture to orchestrate both. Specifically, this architecture **i**) combines Transformers and CNNs in parallel, with Transformers extracting macro information and CNNs capturing micro details, and **ii**) explores across different scales, assessing micro and macro information seamlessly. Additionally, based on the Mesorch architecture, the paper introduces two baseline models aimed at solving IML tasks through mesoscopic representation. Extensive experiments across four datasets have demonstrated that our models surpass the current state-of-the-art in terms of performance, computational complexity, and robustness.

Code — <https://github.com/scu-zjz/Mesorch>

Extended version — <https://arxiv.org/abs/2412.13753>

Introduction

The mesoscopic system exists between macroscopic and microscopic scales. Objects at the mesoscopic scale are large enough to exhibit macroscopic properties, yet they also display interference phenomena related to quantum mechanical phases, similar to microscopic systems. This duality is why it is termed “mesoscopic”—*Yoseph Imry*.

The rapid advancement of multimedia tampering techniques has made detecting and localizing image manipulation more challenging. The ease of creating realistic fakes

*These authors contributed equally.

†Corresponding author: Ji-Zhe Zhou, jzzhou@scu.edu.cn
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

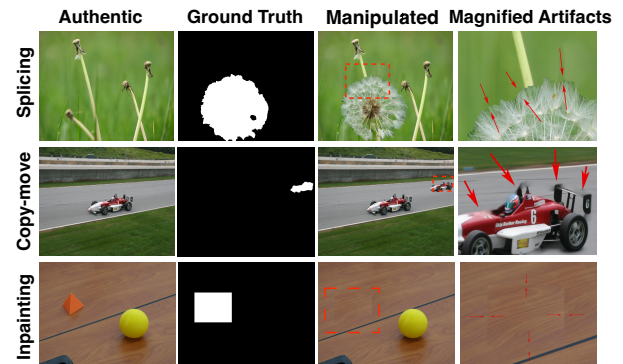


Figure 1: Example of artifacts in three types of tampering. The red dashed box in the third column represents the range of the magnified area in the fourth column. Red arrows in the fourth column point to artifacts that are considered tampering traces.

has fueled tampering incidents and misinformation, highlighting the need for effective forensic methods.

In this context, as illustrated in Figure 1, existing tampering techniques can be broadly categorized into three types (Pun, Yuan, and Bi 2015; Bi, Pun, and Yuan 2016; Wu et al. 2019; Verdoliva 2020; Wei et al. 2023): Splicing (combining parts of different images to create a new one), Copy-move (copying and pasting regions within the same image), and Inpainting (removing and then filling an area with plausible content). Although intricate manual tampering can be imperceptible to the human eye, each type of manipulation still leaves detectable traces at the low level (microscopic level). Therefore, most current techniques in image manipulation localization (IML) consider it a microscopic level task aimed at capturing these tampering traces (artifacts) by extracting microscopic features such as image RGB noise (Zhou et al. 2018; Bayar and Stamm 2018), edge signals (Zhou et al. 2020; Chen et al. 2021), or high-frequency features (Li and Huang 2019; Wang et al. 2022a). These microscopic features are generally effective in revealing artifacts and localizing tampered areas.

However, as depicted in Figure 2, the majority of tampering is typically aimed at deceiving the audience by altering



Figure 2: Random samples from the CASIAv2 dataset. The red line marks the clear boundary of the tampered area. The first column shows tampering that is entirely unrelated to objects, while the other four columns show object-related tampering.

or obscuring the semantics of images. For instance, we observe that in the CASIAv2 dataset, around 80% of the samples are related to the manipulation of objects. Further, the possibility of an object being tampered with will change according to its contribution to the overall semantics of the image. For example, humans and animals in the foreground are more likely to be tampered with than trees and mountains in the background. Therefore, we argue that understanding the object-level (macroscopic-level) semantics is crucial for identifying manipulated regions and delineating suspicious areas. Nonetheless, macroscopic-level semantics alone is insufficient for generating tampering masks, as it lacks necessary details to detect intricate artifacts.

Therefore, figuring out a path to integrate both micro and macro information into the mesoscopic level could be a new solution to improve IML research. DiffForensics (Yu et al. 2024) defined that artifacts exist at the mesoscopic level, which means that they exist between the microscopic and macroscopic levels and share the characteristics of both. However, their research did not delve deeper into how to technically characterize this level. To address this, based on our previous in-depth analysis, we define capturing artifacts at the mesoscopic level by simultaneously capturing microscopic features and extracting macroscopic semantics as a new paradigm for IML tasks. Derived from this paradigm, we introduce the Mesorch (**M**esosopic-**O**rchestration) architecture, which employs parallel encoder and decoder structures specifically designed to represent the mesoscopic level, enabling more effective orchestration between both levels and thus more precisely capturing artifacts at the mesoscopic level.

In the encoder stage, current research (Huang et al. 2016; Zhang et al. 2018; Yuan et al. 2021) has shown that CNN and Transformer models excel in processing microscopic features and macroscopic semantics, respectively. However, most existing IML methods (Wu et al. 2019; Hu et al. 2020; Guillaro et al. 2023) still rely exclusively on either CNNs or

Transformers for decision-making. Although some models, such as ObjectFormer (Wang et al. 2022a), have recognized the advantages of combining CNNs and Transformers, they are designed as sequential models, where the architectures are connected in a linear order. The sequential approach often causes one model to dominate the decision-making process, as most computational resources or parameters are concentrated in that model, overshadowing the strengths of the other. Consequently, the performance does not surpass that of single-architecture models and fails to fully leverage both the microscopic and macroscopic levels. To address this limitation, we adopt a parallel architecture that simultaneously utilizes CNNs and Transformers. This design effectively orchestrates the strengths of both approaches, specifically targeting the capture of artifacts at the mesoscopic level.

Building on this, shallow feature maps provide microscopic features, while deep feature maps offer macroscopic semantics. Thus, using a multi-scale approach in the decoder stage to simultaneously decode feature maps at different scales can help the model more precisely capture artifacts at the mesoscopic level. Similarly, some existing models, such as MVSS-Net (Chen et al. 2021) and Trufor (Guillaro et al. 2023), also utilize multi-scale decoding methods at this stage. However, these models assume equal weighting across all scales without explicitly adjusting the weights, potentially overlooking the differences between scales. This can lead to insufficient utilization of key features or overemphasis on less important features, negatively impacting overall performance. To effectively address this issue, we introduce an adaptive weighting module that dynamically adjusts the importance of each scale. Additionally, by pruning less significant scales, the model significantly reduces computational costs and parameters with minimal impact on performance.

In summary, we propose the Mesorch architecture, a hybrid model combining CNNs and Transformers that dynamically adjusts scale weights to efficiently represent the mesoscopic level. By pruning low-significance scales, we reduce parameters and computational costs, resulting in two baseline models. Testing on four datasets demonstrates SOTA performance in F1 score, robustness, and FLOPs.

Our contributions are threefold:

- We introduce the Mesorch architecture, a hybrid model that leverages the strengths of CNNs and Transformers in parallel. This architecture combines a multi-scale approach to effectively orchestrate microscopic and macroscopic levels, thereby precisely capturing mesoscopic level artifacts in IML tasks.
- We propose an adaptive weighting module that dynamically adjusts the importance of each scale. Additionally, by selectively pruning less impactful scales, our approach considerably reduces computational costs and parameters while maintaining both robustness and performance.
- We develop two baseline models based on the Mesorch architecture. Comprehensive experiments on benchmark datasets demonstrate that our method achieves SOTA performance in F1 score, robustness, and FLOPs.

Related Work

Architectures in Manipulation Localization

CNN-based Models In the realm of image manipulation localization, CNN-based models have long been the mainstream due to their robust feature extraction capabilities, excelling in capturing local textural anomalies indicative of manipulation. Some models like ManTra-Net (Wu et al. 2019) and SPAN (Hu et al. 2020) have been the mainstay in image manipulation localization, utilizing architectures such as VGG (Simonyan and Zisserman 2015) and ResNet-50 (He et al. 2016) to effectively capture local textural anomalies. Recent advancements include integrating contrastive learning, as seen in models like NCL (Zhou et al. 2023), enhancing localization capabilities.

Transformer-based Models Building on the need for broader contextual understanding, Transformer-based models IML-ViT (Ma et al. 2023) is the first to integrate Transformer architectures into the IML domain, TruFor (Guillaro et al. 2023) utilize architectures like SegFormer (Xie et al. 2021). These models are adept at synthesizing wide-ranging contextual information, dynamically focusing on areas of potential manipulation to improve localization accuracy, thus marking a significant advancement in the field. Similarly, MGQFormer (Zeng et al. 2024) uses a query-based Transformer architecture to pinpoint potential manipulation areas. Focusing on contextually relevant features refines manipulation localization precision, illustrating the potent impact of Transformer technologies in delivering robust, context-aware solutions that surpass prior methods.

Hybrid CNN-Transformer Models In addition to the aforementioned models, hybrid approaches that sequentially combine CNNs and Transformers, such as ObjectFormer (Wang et al. 2022a), have also gained significant attention. ObjectFormer employs the EfficientNet architecture as the initial part of its encoder to downsample input data into specific feature blocks, which are then passed sequentially into a dual-stream vision transformer.

Multi-Scale Applications

In the decoder stage, multi-scale techniques have been adopted in image manipulation localization to enhance feature analysis at different resolutions. For example, MVSS-Net (Chen et al. 2021) utilizes a multi-view, multi-scale supervision strategy to detect image manipulations. By incorporating both local edge information and holistic context, MVSS-Net leverages multi-scale feature learning to accurately identify manipulated regions, offering a highly generalizable solution across various datasets and scenarios. Similarly, TruFor (Guillaro et al. 2023) employs the SegFormer (Xie et al. 2021) backbone to effectively integrate spatial relationships across the entire image. Through its multi-scale approach, TruFor enhances the localization of manipulation inconsistencies by combining fine-grained local features with broad contextual information, ensuring robust and accurate manipulation localization.

Method

In this section, we introduce the Mesorch framework, shown in Figure 3. The process begins with an RGB image, which is processed to extract high- and low-frequency features using Discrete Cosine Transform (DCT) (Gonzalez and Woods 2018). These features are then merged with the original image to create the high-frequency and low-frequency enhanced images. The high-frequency enhanced image is passed to the *LocalFeatureModule*, while the low-frequency enhanced image is passed to the *GlobalFeatureModule*. Each module outputs feature maps at four distinct scales. The corresponding decoder processes these feature maps to generate initial predictions of manipulated regions. These multi-scale predictions are weighted and combined to produce the final prediction. To further enhance efficiency, a pruning method is applied after initial model convergence, removing less significant scales to optimize the parameter count and FLOPs.

Integration of Local and Global Information for Enhanced Prediction

Artifacts hidden at the mesoscopic level may not be prominent in the RGB domain, but they can be amplified through frequency domain information (Wang et al. 2022a). Therefore, we extract high-frequency and low-frequency features from the frequency domain to enhance the capabilities of the local and global feature encoder modules.

Feature Enhancement with Discrete Cosine Transform

Initially, as shown in Figure 3, an RGB image \mathbf{x} with size $H \times W \times 3$ is processed using the DCT to separate it into high-frequency components \mathbf{x}_h and low-frequency components \mathbf{x}_l . These components, retaining the dimension $H \times W \times 3$, are then combined with the original image \mathbf{x} to form enhanced representations for further processing:

$$\mathbf{I}_h = \{\mathbf{x}, \mathbf{x}_h\}, \quad \mathbf{I}_h \in \mathbb{R}^{H \times W \times 6} \quad (1)$$

$$\mathbf{I}_l = \{\mathbf{x}, \mathbf{x}_l\}, \quad \mathbf{I}_l \in \mathbb{R}^{H \times W \times 6} \quad (2)$$

Feature Encoding and Scale-Wise Decoding After further enhancing the high-frequency and low-frequency features, the high-frequency enhanced image \mathbf{I}_h is processed by the *LocalFeatureEncoder*, while the low-frequency enhanced image \mathbf{I}_l is processed by the *GlobalFeatureEncoder*. Each encoder subsequently outputs features at four distinct scales:

$$\{L_{s_1}, L_{s_2}, L_{s_3}, L_{s_4}\} = \text{LocalFeatureEncoder}(\mathbf{I}_h), \quad (3)$$
$$L_{s_i} \in \mathbb{R}^{\frac{H}{2^{(i+1)}} \times \frac{W}{2^{(i+1)}} \times C_{i_{\text{local}}}}$$

$$\{G_{s_1}, G_{s_2}, G_{s_3}, G_{s_4}\} = \text{GlobalFeatureEncoder}(\mathbf{I}_l), \quad (4)$$
$$G_{s_i} \in \mathbb{R}^{\frac{H}{2^{(i+1)}} \times \frac{W}{2^{(i+1)}} \times C_{i_{\text{global}}}}$$

Here, $C_{i_{\text{local}}}$ and $C_{i_{\text{global}}}$ denote the total number of output channels at each scale i for the local and global encoders, respectively.

The feature maps generated by the local and global feature encoder at scales $i = 1, 2, 3, 4$ are then processed through

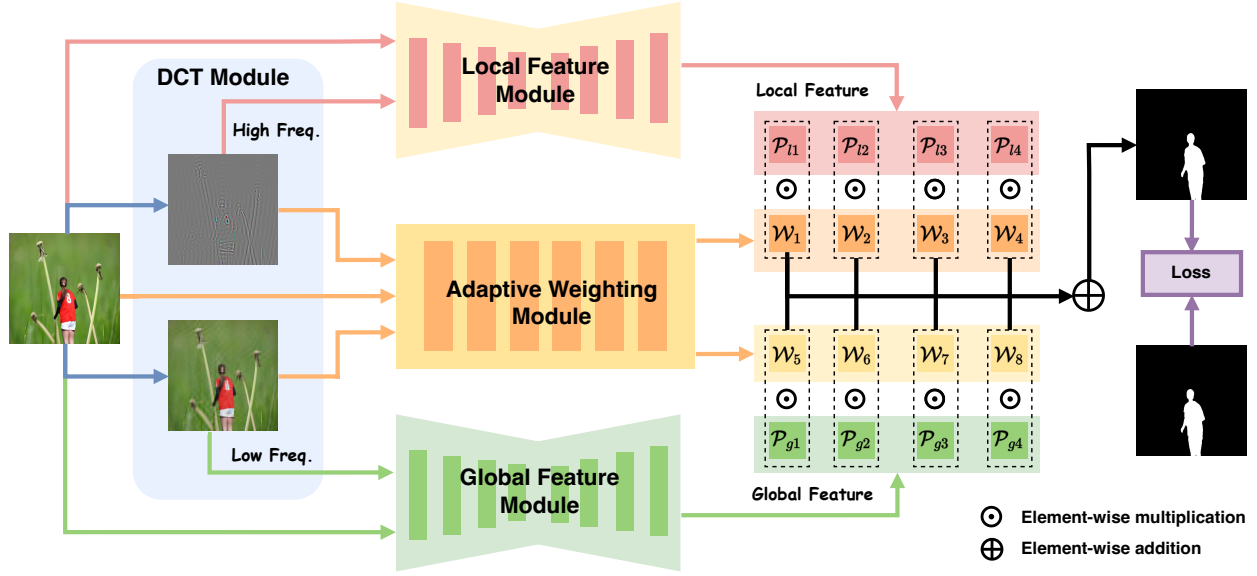


Figure 3: Mesorch Framework: The input RGB image undergoes high- and low-frequency processing in the DCT Module to generate respective high-frequency and low-frequency representations. The Local Feature Module focuses on detecting fine-grained manipulation using both the original and high-frequency images, while the Global Feature Module captures object-level tampering cues by leveraging the original and low-frequency images. The Adaptive Weighting Module dynamically integrates these images by assigning pixel-level weights to local and global features. The final combined features are used for prediction and compared with ground-truth labels to compute the loss.

the decoder, which outputs prediction masks with a shape of $\frac{H}{4} \times \frac{W}{4} \times 1$ for each scale i :

$$P_{l_i} = \text{LocalFeatureDecoder}(L_{s_i}) \quad (5)$$

$$P_{g_i} = \text{GlobalFeatureDecoder}(G_{s_i}) \quad (6)$$

The local and global predictions are combined to produce a summed final prediction mask with a shape of $\frac{H}{4} \times \frac{W}{4} \times 1$. This mask is then resized to the original dimensions of the image ($H \times W$) to generate the final prediction mask P_{final} :

$$P_{\text{final}} = \text{Resize} \left(\sum_{i=1}^4 (P_{l_i} + P_{g_i}), H, W \right) \quad (7)$$

The model’s performance is measured by calculating the cross-entropy loss between P_{final} and the ground truth mask, which identifies the actually manipulated regions in the image. The ground truth mask serves as a binary map highlighting the manipulated areas, guiding the model to focus on discrepancies between the predicted and actual manipulations:

$$\text{Loss} = \text{CrossEntropyLoss}(P_{\text{final}}, \text{Mask}) \quad (8)$$

Adaptive Scale Weighting and Model Pruning

To address the issue of equal weighting across scales, which can lead to inefficient feature utilization, and to reduce the parameter count in the hybrid model, we introduce the AdaptiveWeightingModule and model pruning method to

improve prediction accuracy and computational efficiency.

Scale-Wise Importance The weighting network takes a concatenated input of the original RGB image \mathbf{x} , its high-frequency component, and its low-frequency component, resulting in an input of size $\mathbb{R}^{H \times W \times 9}$. The network then produces a normalized weight vector W , where each element reflects the importance of each scale’s prediction at every pixel:

$$W = \text{WeightingModule}(\{\mathbf{x}, \mathbf{x}_h, \mathbf{x}_l\}), \quad (9)$$

$$W \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times 8}$$

Pixel-Wise Prediction Fusion The final prediction of manipulated regions is determined by performing a weighted sum of the predicted masks across all scales. First, the combined mask P_{all} is formed by merging the local and global predicted masks from each scale, where $P_{\text{all}} \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times 8}$. Then, the weighted summation is calculated:

$$P_{\text{final}} = \text{Resize} \left(\sum_{j=1}^8 W_j \cdot P_{\text{all}_j}, H, W \right) \quad (10)$$

Here, P_{final} is the final prediction mask, resized to the original image dimensions ($H \times W$). As with the earlier prediction, this mask is compared to the ground truth mask to calculate the cross-entropy loss.

Model	F1					Permute F1				
	Coverage	Columbia	NIST16	CASIAv1	Avg.	Coverage	Columbia	NIST16	CASIAv1	Avg.
MVSS-Net	0.4860	0.7399	0.3363	0.5832	0.5364	0.5172	0.7879	0.3775	0.6016	0.5711
PSCC-Net	0.4475	0.8841	0.3457	0.6304	0.5769	0.4930	0.8937	0.3944	0.6382	0.6048
CAT-Net	0.4273	0.9150	0.2521	0.8081	0.6006	0.5165	0.9547	0.3316	0.8154	0.6546
TruFor	0.4573	0.8845	0.3480	0.8176	0.6269	0.5369	0.9547	0.4046	0.8340	0.6826
Mesorch (ours)	0.5862	0.8903	0.3921	<u>0.8398</u>	0.6771	0.6342	0.9708	0.4514	<u>0.8472</u>	0.7259
Mesorch-P (ours)	<u>0.5470</u>	0.9224	<u>0.3888</u>	0.8465	<u>0.6762</u>	0.6107	<u>0.9662</u>	0.4441	0.8564	<u>0.7194</u>

Table 1: Comparison of model performances using standard F1 and permute-F1 metrics, where models denoted with “-P” were trained with pruning methods.

Perturbation	Model	Standard Deviations								Avg.F1
		None	3	7	11	15	19	23		
GaussNoise	MVSS-Net	0.5832	0.5824	0.5822	0.5764	0.5736	0.5620	0.5613	0.5744	
	PSCC-Net	0.6304	0.6127	0.5752	0.5540	0.5402	0.5232	0.5115	0.5639	
	CAT-Net	0.8081	0.7979	0.7883	0.7832	0.7720	0.7573	0.7551	0.7802	
	TruFor	0.8176	0.7666	0.7378	0.7190	0.6947	0.6832	0.6780	0.7286	
	Mesorch(ours)	0.8398	0.8205	0.8050	0.7968	0.7887	0.7780	0.7696	0.7998	
	Mesorch-P(ours)	0.8465	0.8184	0.7872	0.7694	0.7636	0.7543	0.7501	<u>0.7842</u>	
GaussianBlur	Model	Kernel Size								Avg.F1
		None	3	7	11	15	19	23		
		MVSS-Net	0.5832	0.4587	0.3097	0.2369	0.1890	0.1571	0.1392	0.2962
		PSCC-Net	0.6304	0.5410	0.4531	0.3156	0.1655	0.1140	0.0775	0.3282
		CAT-Net	0.8081	0.7512	0.6532	0.5435	0.4337	0.3142	0.2142	0.5312
		TruFor	0.8176	0.7508	0.6881	0.6032	0.4563	0.2741	0.1304	0.5320
	Mesorch(ours)	0.8398	0.7898	0.7081	0.6277	0.5328	0.4193	0.2940	0.6016	
	Mesorch-P(ours)	0.8465	0.7867	0.7179	0.6380	0.5262	0.4106	0.2967	0.6032	
JpegCompression	Model	Quality Factors								Avg.F1
		None	100	90	80	70	60	50		
		MVSS-Net	0.5832	0.5695	0.5446	0.5170	0.4906	0.4489	0.3888	0.5061
		PSCC-Net	0.6304	0.6220	0.5789	0.4930	0.4518	0.3846	0.2869	0.4925
		CAT-Net	0.8081	0.7896	0.7858	0.7431	0.7230	0.6838	0.6132	0.7352
		TruFor	0.8176	0.8060	0.7938	0.7017	0.6852	0.6329	0.4942	0.7049
	Mesorch(ours)	0.8398	0.8312	0.8194	0.7716	0.7706	0.7285	0.6552	0.7738	
	Mesorch-P(ours)	0.8465	0.8314	0.8219	0.7653	0.7598	0.7161	0.6239	<u>0.7664</u>	

Table 2: Robustness test. “Avg.F1” represents the average F1 score on CASIAv1 across perturbation strengths.

Rationale for Secondary Pruning Although the initial training phase allows the model to converge and identify potentially useful features across various scales, further analysis often reveals that some scales might contain redundant or even noisy information, which could diminish the model’s overall effectiveness. Consequently, it becomes essential to evaluate the contribution of each scale using the following criteria:

First, the average weight \bar{W}_i for each scale i is calculated by averaging the weights $W_{i,n}$ across all pixels or feature units n within that scale:

$$\bar{W}_i = \frac{1}{N} \sum_{n=1}^N W_{i,n} \quad (11)$$

Next, the pruning condition is defined as follows:

$$\text{Prune Condition: } \bar{W}_i < \epsilon \quad (12)$$

Here, \bar{W}_i represents the mean weight of the i -th scale, calculated over all N pixels units in that scale. If this average weight falls below a predefined threshold ϵ , the corresponding scale is deemed to contribute minimally to the model and is subsequently pruned. This approach ensures that the model focuses on scales that provide meaningful information while discarding those that add little to the overall performance.

Experiments

Experimental Setup

Training Our model was trained using the standardized Protocol-CAT dataset, provided through a codebase referenced in (Ma et al. 2024). This protocol includes established datasets and typical data augmentation methods. All images were resized to 512x512 pixels. We conducted the training over 150 epochs, utilizing a batch size of 12 on four




Figure 4: Qualitative analysis of SOTA models. We randomly selected and compared four semantically manipulated images and one non-semantically manipulated image based on their respective proportions in the datasets. The first four images are semantically manipulated, while the last one is non-semantically manipulated.

NVIDIA 4090 graphics cards. The learning rate followed a cosine schedule (Loshchilov and Hutter 2017), starting at $1e-4$ and tapering to a minimum of $5e-7$, with a warm-up period of 2 epochs to gradually adjust the learning rate. The AdamW optimizer was used with a weight decay of 0.05 to mitigate overfitting. Furthermore, we set the accumulation iteration to 2, effectively adjusting the batch size to enhance the model’s generalization across diverse data inputs.

Testing We conducted our model evaluations using publicly recognized benchmark, IMDLBenCo (Ma et al. 2024), across four widely used datasets: CASIAv1 (Dong, Wang, and Tan 2013), Coverage (Wen et al. 2016), NIST16 (Guan et al. 2019), and Columbia (Hsu and Chang 2006). These datasets are widely recognized for their diverse challenges and have been instrumental in assessing the generalization of image manipulation localization methods.

Metrics As with our testing, we followed the same publicly recognized benchmarks for evaluation, employing standard pixel-level F1 scores to gauge performance. The results, calculated under the standard threshold of 0.5, provide a comprehensive assessment of localization accuracy.

State-of-the-Art Comparison

To ensure an accurate evaluation, we trained the models using open-source code on the resolutions recommended in their respective papers, employing the CAT-Net (Kwon et al.

2022) protocol dataset. We then benchmarked their performance across several well-established datasets using the F1 score. The comparative analysis included various methods such as PSCC-Net (Liu et al. 2022a), MVSS-Net (Chen et al. 2021), CAT-Net (Kwon et al. 2022), and Trufor (Guillaro et al. 2023).

Localization Results In Table 1, we have highlighted the best-performing model in bold and the second-best model with an underline across all evaluated datasets. Our proposed method, both before and after pruning, consistently ranks among the top performers, achieving either the highest or second-highest results. This consistent performance demonstrates the accuracy of our approach in effectively handling various image manipulation localization tasks. Additionally, Figure 4 qualitatively shows that our model successfully captures both the object layout and fine details at the mesoscopic level, resulting in highly accurate manipulated masks.

Robust Performance To assess model robustness under various conditions, we tested on the CASIAv1 dataset and reported the results in Table 2. We introduced Gaussian noise with varying standard deviations, Gaussian blur with different kernel sizes, and JPEG compression using various quality factors as perturbations. The results show that our model consistently achieved SOTA robustness across all three perturbations. Notably, even after pruning, our method

Model	Parameters (M)	FLOPs (G)
MVSS	150.528	171.008
PSCC	3.668	376.832
CAT-Net	116.736	137.216
TruFor	68.697	236.544
Mesorch(ours)	85.754	124.928
Mesorch-P(ours)	62.235	64.821

Table 3: Comparison of parameters and computational efficiency (Flops) across different models.

Model	Multi.	Weight	Extractor	Prune	Avg.F1
CNN	×	×	×	×	0.5610
	✓	×	×	×	0.6031
Transformer	×	×	×	×	0.5723
	✓	×	×	×	0.6096
CNN+Trans.	×	×	×	×	0.6422
	✓	×	×	×	0.6501
	✓	×	✓	×	0.6612
	✓	✓	×	×	0.6653
	✓	✓	✓	×	0.6771
	✓	–	✓	✓	0.6762

Table 4: Independent Evaluation. “Avg.F1” represents the average F1 score across four datasets.

maintained superior robustness compared to all previous models, demonstrating its effectiveness in handling diverse image distortions.

FLOPs and Parameters The number of parameters and FLOPs for all measurements were calculated based on a resolution of 512x512 and a batch size of 1. As shown in Table 3, our results demonstrate that our model has fewer FLOPs than all SOTA models, with a parameter count second only to PSCC-Net. Furthermore, the application of our pruning method further reduces FLOPs and total parameter count, giving our model a greater advantage over existing SOTA models.

CNN	Transformer	Avg.F1
ConvNeXt	MAE	0.6215
	PVT	0.6704
	Segformer	0.6771
	Swin	0.6089
Resnet	MAE	0.6327
	PVT	0.6496
	Segformer	0.6615
	Swin	0.6022

Table 5: The performance of the same module uses different backbones. “Avg.F1” stands for “Average F1,” representing the mean value of the standard F1 score across four datasets.

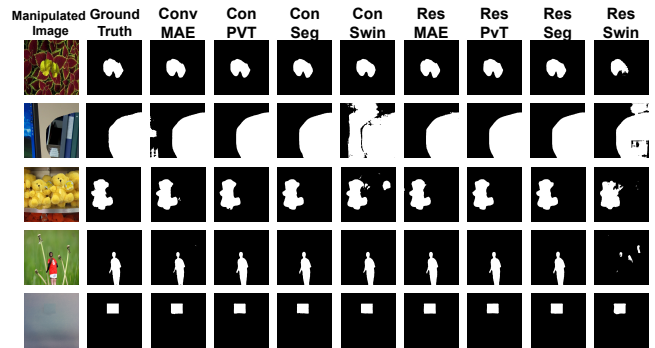


Figure 5: Ablation study on different backbones implemented in Mesorch. Qualitatively “Conv” is ConvNeXt, “Res” is ResNet, and “Seg” is Segformer.

Ablation Study

Independent Evaluation of Proposed Method To validate the effectiveness of our proposed method, we first independently evaluated the ConvNeXt architecture for CNNs and the Segformer architecture for Transformers as baseline models and assessed their performance under multi-scale conditions. After establishing the baseline performance, we further evaluated the hybrid model, gradually incorporating the multi-scale approach, weighting method, and DCT feature extraction. We also tested the pruned model, as shown in Table 4. The results demonstrate that each component proposed in the paper is crucial for accurately localizing image manipulations.

Comparative Analysis of Model Architectures We evaluated the performance of the Mesorch architecture using different combinations of two CNN models (Resnet-50 (He et al. 2016) and ConvNeXt-Tiny (Liu et al. 2022b)) with four Transformer models (MAE-Base (He et al. 2022), PVT-B3 (Wang et al. 2022b), Segformer-B3 (Xie et al. 2021), and SwinTransformer-Base (Liu et al. 2021)). The performance differences are summarized in 5. Additionally, Figure 5 provides qualitative analysis, showcasing how the Mesorch architecture performs with different backbones. Our findings indicate that the combination of ConvNeXt and Segformer excels in both macroscopic localization and capturing microscopic features, outperforming other model combinations.

Conclusion

Inspired by the mesoscopic perspective, this paper redefines the IML task to orchestrate both microscopic and macroscopic levels. Building on this, we propose the Mesorch architecture, a hybrid model that leverages the strengths of CNNs and Transformers while dynamically adjusting scale weights to efficiently capture artifacts at the mesoscopic level. To reduce the parameter count and computational cost, we also introduce two baseline models based on this architecture. Extensive testing on large-scale datasets demonstrates that our approach consistently achieves SOTA performance in terms of F1 score, robustness, and FLOPs.

Acknowledgments

The authors would like to express our gratitude to Professor Bin Li for introducing the mesoscopic concept in Diff-Forensics, which inspired our work. We also thank Xianhang Cheng and Kaiwen Feng for their contributions to the design of Figure 3.

This work was jointly supported by the Sichuan Natural Science Foundation under grant 2024YFHZ0355, Sichuan Major Projects under grant 2024ZDZX0001, and the Science and Technology Development Fund, Macau SAR, under grants 0141/2023/RIA2 and 0193/2023/RIA3. Numerical computations were jointly supported by Hefei Advanced Computing Center and Chengdu Haiguang Integrated Circuit Design Co., ltd. with HYGON K100AI DCU units.

References

- Bayar, B.; and Stamm, M. C. 2018. Constrained Convolutional Neural Networks: A New Approach Towards General Purpose Image Manipulation Detection. *IEEE Transactions on Information Forensics and Security*, 13(11): 2691–2706.
- Bi, X.; Pun, C.-M.; and Yuan, X.-C. 2016. Multi-level dense descriptor and hierarchical feature matching for copy-move forgery detection. *Information Sciences*, 345: 226–242.
- Chen, X.; Dong, C.; Ji, J.; Cao, J.; and Li, X. 2021. Image Manipulation Detection by Multi-View Multi-Scale Supervision. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 14165–14173. Montreal, QC, Canada: IEEE. ISBN 978-1-66542-812-5.
- Dong, J.; Wang, W.; and Tan, T. 2013. CASIA Image Tampering Detection Evaluation Database. In *2013 IEEE China Summit and International Conference on Signal and Information Processing*, 422–426. Beijing, China: IEEE. ISBN 978-1-4799-1043-4.
- Gonzalez, R. C.; and Woods, R. E. 2018. *Digital Image Processing*. Pearson, 4th edition.
- 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 *2019 IEEE Winter Applications of Computer Vision Workshops (WACVW)*, 63–72. Waikoloa Village, HI, USA: IEEE. ISBN 978-1-72811-392-0.
- 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.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked Autoencoders Are Scalable Vision Learners. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 15979–15988. New Orleans, LA, USA: IEEE. ISBN 978-1-66546-946-3.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778. Las Vegas, NV, USA: IEEE. ISBN 978-1-4673-8851-1.
- Hsu, Y.-f.; and Chang, S.-f. 2006. Detecting Image Splicing using Geometry Invariants and Camera Characteristics Consistency. In *2006 IEEE International Conference on Multimedia and Expo*, 549–552. Toronto, ON, Canada: IEEE. ISBN 978-1-4244-0367-7.
- Hu, X.; Zhang, Z.; Jiang, Z.; Chaudhuri, S.; Yang, Z.; and Nevatia, R. 2020. SPAN: Spatial pyramid attention network for image manipulation localization. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, 312–328. Springer.
- Huang, S.; Xu, Z.; Tao, D.; and Zhang, Y. 2016. Part-stacked CNN for fine-grained visual categorization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1173–1182.
- 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, H.; and Huang, J. 2019. Localization of Deep Inpainting Using High-Pass Fully Convolutional Network. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 8300–8309. Seoul, Korea (South): IEEE. ISBN 978-1-72814-803-8.
- Liu, X.; Liu, Y.; Chen, J.; and Liu, X. 2022a. Psc-net: Progressive spatio-channel correlation network for image manipulation detection and localization. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(11): 7505–7517.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 9992–10002. Montreal, QC, Canada: IEEE. ISBN 978-1-66542-812-5.
- Liu, Z.; Mao, H.; Wu, C.-Y.; Feichtenhofer, C.; Darrell, T.; and Xie, S. 2022b. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11976–11986.
- Loshchilov, I.; and Hutter, F. 2017. SGDR: Stochastic Gradient Descent with Warm Restarts. (arXiv:1608.03983). ArXiv:1608.03983 [cs, math].
- Ma, X.; Du, B.; Liu, X.; Hammadi, A. Y. A.; and Zhou, J. 2023. Iml-vit: Image manipulation localization by vision transformer.
- Ma, X.; Zhu, X.; Su, L.; Du, B.; Jiang, Z.; Tong, B.; Lei, Z.; Yang, X.; Pun, C.-M.; Lv, J.; et al. 2024. IMDL-BenCo: A Comprehensive Benchmark and Codebase for Image Manipulation Detection & Localization.
- Pun, C.-M.; Yuan, X.-C.; and Bi, X.-L. 2015. Image forgery detection using adaptive oversegmentation and feature point matching. *IEEE transactions on information forensics and security*, 10(8): 1705–1716.
- Simonyan, K.; and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. (arXiv:1409.1556). ArXiv:1409.1556 [cs].

- Verdoliva, L. 2020. Media Forensics and DeepFakes: An Overview. *IEEE Journal of Selected Topics in Signal Processing*, 14(5): 910–932.
- Wang, J.; Wu, Z.; Chen, J.; Han, X.; Shrivastava, A.; Lim, S.-N.; and Jiang, Y.-G. 2022a. ObjectFormer for Image Manipulation Detection and Localization. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2354–2363. New Orleans, LA, USA: IEEE. ISBN 978-1-66546-946-3.
- Wang, W.; Xie, E.; Li, X.; Fan, D.-P.; Song, K.; Liang, D.; Lu, T.; Luo, P.; and Shao, L. 2022b. Pvt v2: Improved baselines with pyramid vision transformer. *Computational Visual Media*, 8(3): 415–424.
- Wei, Y.; Xiao, B.; Bi, X.; Ma, Z.; Liu, Y.; and Ma, Z. 2023. Secondary Labeling: A Novel Labeling Strategy for Image Manipulation Detection. In *Proceedings of the 31st ACM International Conference on Multimedia*, 8225–8232.
- 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 *2016 IEEE International Conference on Image Processing (ICIP)*, 161–165. Phoenix, AZ, USA: IEEE. ISBN 978-1-4673-9961-6.
- Wu, Y.; et al. 2019. ManTra-Net: Manipulation Tracing Network for Detection and Localization of Image Forgeries With Anomalous Features. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9535–9544. Long Beach, CA, USA: IEEE. ISBN 978-1-72813-293-8.
- Xie, E.; Wang, W.; Yu, Z.; Anandkumar, A.; Alvarez, J. M.; and Luo, P. 2021. SegFormer: Simple and efficient design for semantic segmentation with transformers. *Advances in neural information processing systems*, 34: 12077–12090.
- Yu, Z.; Ni, J.; Lin, Y.; Deng, H.; and Li, B. 2024. DiffForensics: Leveraging Diffusion Prior to Image Forgery Detection and Localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12765–12774.
- Yuan, Y.; Fu, R.; Huang, L.; Lin, W.; Zhang, C.; Chen, X.; and Wang, J. 2021. Hrformer: High-resolution transformer for dense prediction.
- 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.
- Zhang, Q.; Cao, R.; Shi, F.; Wu, Y. N.; and Zhu, S.-C. 2018. Interpreting CNN knowledge via an explanatory graph. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Zhou, J.; Ma, X.; Du, X.; Alhammadi, A. Y.; and Feng, W. 2023. Pre-training-free Image Manipulation Localization through Non-Mutually Exclusive Contrastive Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 22346–22356.
- 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. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(07): 13058–13065.
- Zhou, P.; Han, X.; Morariu, V. I.; and Davis, L. S. 2018. Learning Rich Features for Image Manipulation Detection. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1053–1061. Salt Lake City, UT, USA: IEEE. ISBN 978-1-5386-6420-9.