

Image Manipulation Detection by Multi-View Multi-Scale Supervision

Xinru Chen^{1,2*}, Chengbo Dong^{1,2*}, Jiaqi Ji^{1,2}, Juan Cao^{3,4}, Xirong Li^{1,2†}

¹MoE Key Lab of Data Engineering and Knowledge Engineering, Renmin University of China

²AIMC Lab, School of Information, Renmin University of China

³Institute of Computing Technology, Chinese Academy of Sciences

⁴State Key Laboratory of Media Convergence Production Technology and Systems

<https://github.com/dong03/MVSS-Net>

Abstract

The key challenge of image manipulation detection is how to learn generalizable features that are sensitive to manipulations in novel data, whilst specific to prevent false alarms on authentic images. Current research emphasizes the sensitivity, with the specificity overlooked. In this paper we address both aspects by multi-view feature learning and multi-scale supervision. By exploiting noise distribution and boundary artifact surrounding tampered regions, the former aims to learn semantic-agnostic and thus more generalizable features. The latter allows us to learn from authentic images which are nontrivial to be taken into account by current semantic segmentation network based methods. Our thoughts are realized by a new network which we term MVSS-Net. Extensive experiments on five benchmark sets justify the viability of MVSS-Net for both pixel-level and image-level manipulation detection.

1. Introduction

Digital images can now be manipulated with ease and often in a visually imperceptible manner [11]. *Copy-move* (copy and move elements from one region to another region in a given image), *splicing* (copy elements from one image and paste them on another image) and *inpainting* (removal of unwanted elements) are three common types of image manipulation that could lead to misinterpretation of the visual content [1, 19, 23]. This paper targets at auto-detection of images subjected to these types of manipulation. We aim to not only discriminate manipulated images from the authentic, but also pinpoint tampered regions at the pixel level.

Unsurprisingly, the state-of-the-arts are deep learning based [14, 21, 26, 27, 29], specifically focusing on pixel-level manipulation detection [21, 26, 29]. With only two

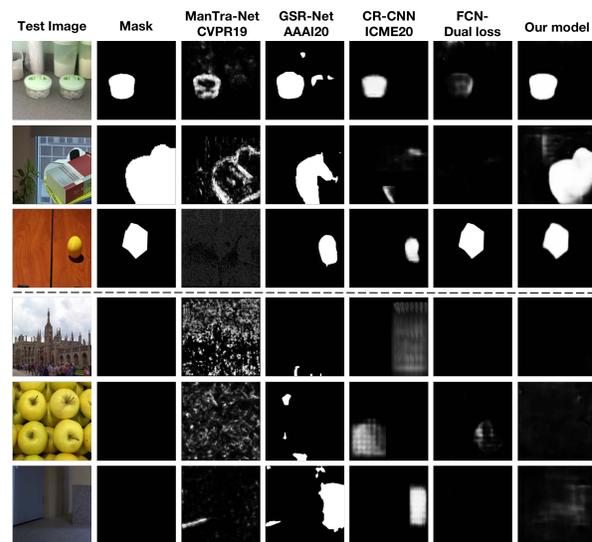


Figure 1. **Image manipulation detection by the state-of-the-arts.** The first three rows are copy-move, splicing and inpainting, followed by three authentic images (thus with blank mask). Our model strikes a good balance between sensitivity and specificity.

classes (*manipulated* versus *authentic*) in consideration, the task appears to be a simplified case of image semantic segmentation. However, an off-the-shelf semantic segmentation network is suboptimal for the task, as it is designed to capture semantic information, making the network dataset-dependent and do not generalize. Prior research [29] reports that DeepLabv2 [4] trained on the CASIAv2 dataset [8] performs well on the CAISAv1 dataset [7] homologous to CASIAv2, yet performs poorly on the non-homologous COVER dataset [25]. A similar behavior of FCN [18] is also observed in this study. Hence, the key question is how to design and train a deep neural network capable of learning *semantic-agnostic* features that are *sensitive* to manipulations, whilst *specific* to prevent false alarms?

In order to learn semantic-agnostic features, image con-

*Xinru Chen and Chengbo Dong contribute equally to this work.

†Corresponding author: Xirong Li (xirong@ruc.edu.cn).

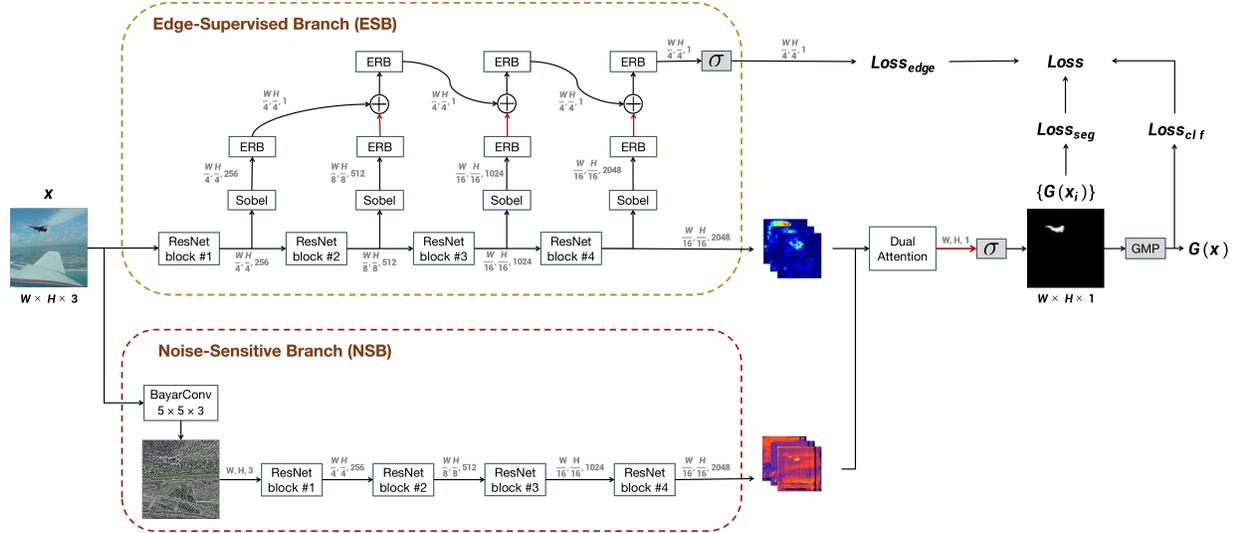


Figure 2. **Conceptual diagram of the proposed *MVSS-Net* model.** We use the edge-supervised branch and the noise-sensitive branch to learn semantic-agnostic features for manipulation detection, and multi-scale supervision to strike a balance between model sensitivity and specificity. Non-trainable layers such as sigmoid (σ) and global max pooling (GMP) are shown in gray.

ment has to be suppressed. Depending on at what stage the suppression occurs, we categorize existing methods into two groups, *i.e.* noise-view methods [14, 16, 26, 27, 30] and edge-supervised methods [21, 29]. Given the hypothesis that novel elements introduced by slicing and/or inpainting differ from the authentic part in terms of their noise distributions, the first group of methods aim to exploit such discrepancy. The noise map of an input image, generated either by pre-defined high-pass filters [9] or by their trainable counterparts [2, 16], is fed into a deep network, either alone [16, 27] or together with the input image [14, 26, 30]. Note that the methods are ineffective for detecting copy-move which introduces no new element. The second group of methods concentrate on finding boundary artifact as manipulation trace around a tampered region, implemented by adding an auxiliary branch to reconstruct the region’s edge [21, 29]. Note that the prior art [29] uniformly concatenates features from different layers of the backbone as input of the auxiliary branch. As such, there is a risk that deeper-layer features, which are responsible for manipulation detection, remain semantic-aware and thus not generalizable.

To measure a model’s generalizability, a common evaluation protocol [14, 21, 26, 29] is to first train the model on a public dataset, say CASIAv2 [8], and then test it on other public datasets such as NIST16 [12], Columbia [13], and CASIAv1 [7]. To our surprise, however, the evaluation is performed exclusively on manipulated images, with metrics w.r.t pixel-level manipulation detection reported. The specificity of the model, which reveals how it handles authentic images and is thus crucial for real-world usability, is ignored. As shown in Fig. 1, their serious false alarm over

authentic images leads to unavailability in practical work. In fact, as current methods [14, 21, 26] mainly use pixel-wise segmentation losses to which an authentic example can contribute is marginal, it is nontrivial for these methods to improve their specificity by learning from the authentic.

Inspired by the Border Network [28], which aggregates features progressively to predict object boundaries, and LesionNet [24] that incorporates an image classification loss for retinal lesion segmentation, we propose *multi-view* feature learning with *multi-scale* supervision for image manipulation detection. To the best of our knowledge (Table 1), we are the first to jointly exploit the noise view and the boundary artifact to learn manipulation detection features. Moreover, such a joint exploitation is nontrivial. To combine the best of the two worlds, new network structures are needed. Our contributions are as follows:

- We propose *MVSS-Net* as a new network for image manipulation detection. As shown in Fig. 2, *MVSS-Net* contains novel elements designed for learning semantic-agnostic and thus more generalizable features.
- We train *MVSS-Net* with multi-scale supervision, allowing us to learn from authentic images, which are ignored by the prior art, and consequently improve the model specificity substantially.
- Extensive experiments on two training sets and five test sets show that *MVSS-Net* compares favorably against the state-of-the-art.

2. Related Work

This paper is inspired by a number of recent works that made novel attempts to learn semantic-agnostic features for

Methods	Views			Backbone	Scales of Supervision		
	RGB	Noise	Fusion		pixel	edge	image
Bappy <i>et al.</i> 2017, J-LSTM [1]	+	-	-	Patch-LSTM	+	-	-
Salloum <i>et al.</i> 2017, MFCN [21]	+	-	-	FCN	+	+	-
Zhou <i>et al.</i> 2020, GSR-Net [29]	+	-	-	Deeplabv2	+	+	-
Li & Huang 2019, HP-FCN [16]	-	High-pass filters	-	FCN	+	-	-
Yang <i>et al.</i> 2020, CR-CNN [27]	-	BayarConv2D	-	Mask R-CNN	+	-	-
Zhou <i>et al.</i> 2018, RGB-N [30]	+	SRM filter	late fusion (bilinear pooling)	Faster R-CNN	*	-	-
Wu <i>et al.</i> 2019, ManTra-Net [26]	+	SRM filter, BayarConv2D	early fusion (feature concatenation)	Wider VGG	+	-	-
Hu <i>et al.</i> 2020, SPAN [14]	+	SRM filter BayarConv2D	early fusion (feature concatenation)	Wider VGG	+	-	-
<i>MVSS-Net(This paper)</i>	+	BayarConv2D	late fusion (dual attention)	FCN	+	+	+

Table 1. **A taxonomy of the state-of-the-art for image manipulation detection.** Note that edge and image labels used in this work are automatically extracted from pixel-level annotations. So our multi-scale supervision does not use extra manual annotation.

image manipulation detection, see Table 1. In what follows, we describe in brief how these attempts are implemented and explain our novelties accordingly. We focus on deep learning approaches to copy-move / splicing / inpainting detection. For the detection of low-level manipulations such as Gaussian Blur and JPEG compression, we refer to [2].

In order to suppress the content information, Li and Huang [16] propose to implement an FCN’s first convolutional layer with trainable high-pass filters and apply their HP-FCN for inpainting detection. Yang *et al.* use BayarConv as the initial convolution layer of their CR-CNN [27]. Although such constrained convolutional layers are helpful for extracting noise information, using them alone brings in the risk of losing other useful information in the original RGB input. Hence, we see an increasing number of works on exploiting information from both the RGB view and the noise view [14, 26, 30]. Zhou *et al.* [30] develop a two-stream Faster R-CNN, coined RGB-N, which takes as input the RGB image and its noise counterpart generated by the SRM filter [9]. Wu *et al.* [26] and Hu *et al.* [14] use both BayarConv and SRM. Given features from distinct views, the need for feature fusion is on. Feature concatenation at an early stage is adopted by [14, 26]. Our *MVSS-Net* is more close to RGB-N as both perform feature fusion at the late stage. However, different from the non-trainable bilinear pooling used in RGB-N, Dual Attention used in *MVSS-Net* is trainable and thus more selective.

As manipulating a specific region in a given image inevitably leaves traces between the tampered region and its surrounding, how to exploit such edge artifact also matters for manipulation detection. Salloum *et al.* develop a multi-task FCN to symmetrically predict a tampered area and its boundary [21]. In a more recent work [29], Zhou *et al.* introduce an edge detection and refinement branch which accepts features from different levels. Given that region segmentation and edge detection are intrinsically two distinct tasks, the challenge lies in how to strike a proper balance

between the two. Directly using deeper features for edge detection as done in [21] has the risk of affecting the main task of manipulation segmentation, while putting all features together as used in [29] may let the deeper features be ignored by the edge branch. Our *MVSS-Net* has an edge-supervised branch that effectively resolves these issues.

Last but not least, we observe that the specificity of an image manipulation detector, *i.e.* how it responds to authentic images, is seldom reported. In fact, the mainstream solutions are developed within an image semantic segmentation network. Naturally, they are trained and also evaluated on manipulated images in the context of manipulation segmentation [29]. The absence of authentic images both in the training and test stages naturally raises concerns regarding the specificity of the detector. In this paper we make a novel attempt to include authentic images for training and test, an important step towards real-world deployment.

3. Proposed Model

Given an RGB image x of size $W \times H \times 3$, we aim for a multi-head deep network G that not only determines whether the image has been manipulated, but also pinpoints its manipulated pixels. Let $G(x)$ be the network-estimated probability of the image being manipulated. In a similar manner we define $G(x_i)$ as pixel-wise probability, with $i = 1, \dots, W \times H$. Accordingly, we denote a full-size segmentation map as $\{G(x_i)\}$. As the image-level decision is naturally subject to pixel-level evidence, we obtain $G(x)$ by Global Max Pooling (GMP) over the segmentation map, *i.e.*

$$G(x) \leftarrow \text{GMP}(\{G(x_i)\}). \quad (1)$$

In order to extract generalizable manipulation detection features, we present a new network that accepts both RGB and noise views of the input image. To strike a proper balance between detection sensitivity and specificity, the multi-

view feature learning process is jointly supervised by annotations of three scales, *i.e.* pixel, edge and image.

3.1. Multi-View Feature Learning

As shown in Fig. 2, *MVSS-Net* consists of two branches, with ResNet-50 as their backbones. The edge-supervised branch (ESB) at the top is specifically designed to exploit subtle boundary artifact around tampered regions, whilst the noise-sensitive branch (NSB) at the bottom aims to capture the inconsistency between tampered and authentic regions. Both clues are meant to be semantic-agnostic.

3.1.1 Edge-Supervised Branch

Ideally, with edge supervision, we hope the response area of the network will be more concentrated on tampered regions. Designing such an edge-supervised network is non-trivial. As noted in Section 2, the main challenge is how to construct an appropriate input for the edge detection head. On one hand, directly using features from the last ResNet block is problematic, as this will enforce the deep features to capture low-level edge patterns and consequently affect the main task of manipulation segmentation. While on the other hand, using features from the initial blocks is also questionable, as subtle edge patterns contained in these shallow features can vanish with ease after multiple deep convolutions. A joint use of both shallow and deep features is thus necessary. However, we argue that simple feature concatenation as previously used in [29] is suboptimal, as the features are mixed and there is no guarantee that the deeper features will receive adequate supervision from the edge head. To conquer the challenge, we propose to construct the input of the edge head in a shallow-to-deep manner.

As illustrated in Fig. 2, features from different ResNet blocks are combined in a progressive manner for manipulation edge detection. In order to enhance edge-related patterns, we introduce a Sobel layer, see Fig. 3(a). Features from the i -th block first go through the Sobel layer followed by an edge residual block (ERB), see Fig. 3(b), before they are combined (by summation) with their counterparts from the next block. To prevent the effect of accumulation, the combined features go through another ERB (top in Fig. 2) before the next round of feature combination. We believe such a mechanism helps prevent extreme cases in which deeper features are over-supervised or fully ignored by the edge head. By visualizing feature maps of the last ResNet block in Fig. 4, we observe that the proposed ESB indeed produces a more focused response near tampered regions.

The output of ESB has two parts: feature maps from the last ResNet block, denoted as $\{f_{esb,1}, \dots, f_{esb,k}\}$, to be used for the main task, and the predicted manipulation edge map, denoted as $\{G_{edge}(x_i)\}$, obtained by transforming the output of the last ERB with a sigmoid (σ) layer. The

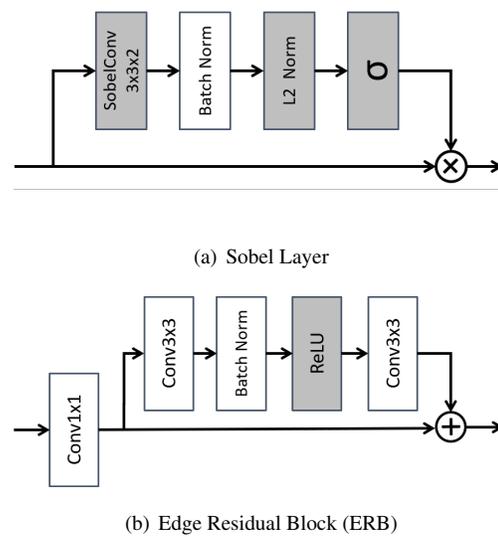


Figure 3. **Diagrams of (a) Sobel layer and (b) edge residual block**, used in ESB for manipulation edge detection.

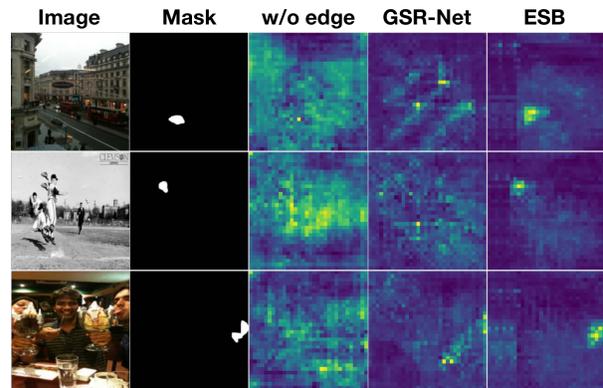


Figure 4. **Visualization of averaged feature maps of the last ResNet block**, brighter color indicating a higher response. Manipulation from the top to bottom is inpainting, copy-move and splicing. Read from the third column are *w/o edge*, *i.e.* ResNet without any edge residual block, *GSR-Net*, *i.e.* ResNet with the GSR-Net alike edge branch, and the proposed *ESB*, which produces a more focused response near tampered regions.

data-flow of this branch is conceptually expressed by Eq. 2,

$$\left. \begin{matrix} [f_{esb,1}, \dots, f_{esb,k}] \\ \{G_{edge}(x_i)\} \end{matrix} \right\} \leftarrow \text{ERB-ResNet}(x). \quad (2)$$

3.1.2 Noise-Sensitive Branch

In order to fully exploit the noise view, we build a noise-sensitive branch (NSB) parallel to ESB. NSB is implemented as a standard FCN (another ResNet-50 as its backbone). Regarding the choice of noise extraction, we adopt BayarConv [2], which is found to be better than the SRM

filter [27]. The output of this branch is an array of k feature maps from the last ResNet block of its backbone, *i.e.*

$$\{f_{nsb,1}, \dots, f_{nsb,k}\} \leftarrow \text{ResNet}(\text{BayarConv}(x)). \quad (3)$$

3.1.3 Branch Fusion by Dual Attention

Given two arrays of feature maps $\{f_{esb,1}, \dots, f_{esb,k}\}$ and $\{f_{nsb,1}, \dots, f_{nsb,k}\}$ from ESB and NSB, we propose to fuse them by a trainable Dual Attention (DA) module [10]. This is new, because previous work [30] uses bilinear pooling for feature fusion, which is non-trainable.

The DA module has two attention mechanisms working in parallel: channel attention (CA) and position attention (PA), see Fig. 5. CA associates channel-wise features to selectively emphasize interdependent channel feature maps. Meanwhile, PA selectively updates features at each position by a weighted sum of the features at all positions. The outputs of CA and PA are summed up, and transformed into a feature map of size $\frac{W}{16} \times \frac{H}{16}$, denoted as $\{G'(x_i)\}$, by a 1×1 convolution. With parameter-free bilinear upsampling followed by an element-wise sigmoid function, $\{G'(x_i)\}$ is transformed into the final segmentation map $\{G(x_i)\}$. Fusion by dual attention is conceptually expressed as

$$\begin{cases} \{G'(x_i)\} \leftarrow DA([f_{esb,1}, \dots, f_{esb,k}, f_{nsb,1}, \dots, f_{nsb,k}]), \\ \{G(x_i)\} \leftarrow \sigma(\text{bilinear-upsampling}(\{G'(x_i)\})). \end{cases} \quad (4)$$

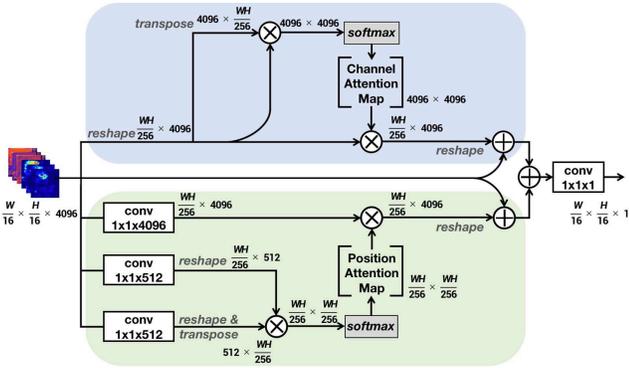


Figure 5. **Dual Attention**, with its channel attention module shown in blue and its position attention module shown in green.

3.2. Multi-Scale Supervision

We consider losses at three scales, each with its own target, *i.e.* a pixel-scale loss for improving the model’s sensitivity for pixel-level manipulation detection, an edge loss for learning semantic-agnostic features and an image-scale loss for improving the model’s specificity for image-level manipulation detection.

Pixel-scale loss. As manipulated pixels are typically in minority in a given image, we use the Dice loss, found to be

effective for learning from extremely imbalanced data [24]:

$$\text{loss}_{seg}(x) = 1 - \frac{2 \cdot \sum_{i=1}^{W \times H} G(x_i) \cdot y_i}{\sum_{i=1}^{W \times H} G^2(x_i) + \sum_{i=1}^{W \times H} y_i^2}, \quad (5)$$

where $y_i \in \{0, 1\}$ is a binary label indicating whether the i -th pixel is manipulated.

Edge loss. As pixels of an edge are overwhelmed by non-edge pixels, we again use the Dice loss for manipulation edge detection, denoted as loss_{edg} . Since manipulation edge detection is an auxiliary task, we do not compute the loss_{edg} at the full size of $W \times H$. Instead, the loss is computed at a smaller size of $\frac{W}{4} \times \frac{H}{4}$, see Fig. 2. This strategy reduces computational cost during training, and in the meanwhile, improves the performance slightly.

Image-scale loss. In order to reduce false alarms, authentic images have to be taken into account in the training stage. This is however nontrivial for the current works [16, 21, 26, 29] as they all rely on segmentation losses. Consider the widely used binary cross-entropy (BCE) loss for instance. An authentic image with a small percent of its pixels misclassified contributes marginally to the BCE loss, making it difficult to effectively reduce false alarms. Also note that the Dice loss cannot handle the authentic image by definition. Therefore, an image-scale loss is needed. We adopt the image-scale BCE loss:

$$\text{loss}_{clf}(x) = -(y \cdot \log G(x) + (1 - y) \cdot \log(1 - G(x))) \quad (6)$$

where $y = \max(\{y_i\})$.

Combined loss. We use a convex combination of the three losses:

$$\text{Loss} = \alpha \cdot \text{loss}_{seg} + \beta \cdot \text{loss}_{clf} + (1 - \alpha - \beta) \cdot \text{loss}_{edg} \quad (7)$$

where $\alpha, \beta \in (0, 1)$ are weights. Note that authentic images are only used to compute loss_{clf} .

4. Experiments

4.1. Experimental Setup

Datasets. For the ease of a head-to-head comparison with the state-of-the-art, we adopt CASIAv2 [8] for training and COVER [25], Columbia [13], NIST16 [12] and CASIAv1 [7] for testing. Meanwhile, we notice DEFACTO [19], a recently released large-scale dataset, containing 149k images sampled from MS-COCO [17] and auto-manipulated by copy-move, splicing and inpainting. Considering the challenging nature of DEFACTO, we choose to perform our ablation study on this new set. As the set has no authentic images, we construct a training set termed DEFACTO-84k, by randomly sampling 64k positive images from DEFACTO and 20k negative images from MS-COCO. In a similar manner, we build a test set termed DEFACTO-12k, by randomly sampling 6k positive images from the

remaining part of DEFACTO and 6k negatives from MSCOCO. Note that to avoid any data leakage, for manipulated images used for training (test), their source images are not included in the test (training) set. In total, our experiments use two training sets and five test sets, see Table 2.

Dataset	Negative	Positive	cpmv	spli	inpa
<i>Training</i>					
DEFACTO-84k [19]	20,000	64,417	12,777	34,133	17,507
CASIAv2 [8]	7,491	5,063	3,235	1,828	0
<i>Test</i>					
COVER [25]	100	100	100	0	0
Columbia [13]	183	180	0	180	0
NISTI6 [12]	0	564	68	288	208
CASIAv1 [7]	800	920	459	461	0
DEFACTO-12k [19]	6,000	6,000	2,000	2,000	2,000

Table 2. **Two training sets and five test sets in our experiments.** DEFACTO-84k and DEFACTO-12k are used for training and test in the ablation study (Section 4.2), while for the SOTA comparison (Section 4.3) we train on CASIAv2 and evaluate on all test sets.

Evaluation Criteria. For pixel-level manipulation detection, following previous works [21, 29, 30], we compute pixel-level precision and recall, and report their F1. For image-level manipulation detection, in order to measure the miss detection rate and false alarm rate, we report sensitivity, specificity and their F1. AUC, as a decision-threshold-free metric, is also reported. Authentic images per test set are only used for image-level evaluation. For both pixel-level and image-level F1 computation, the default threshold is 0.5, unless otherwise stated.

The overall performance is measured by Com-F1, defined as the harmonic mean of pixel-level and image-level F1. Com-F1 is sensitive to the lowest value of pixel-F1 and image-F1. In particular, it scores 0 when either pixel-F1 or image-F1 is 0, which does not hold for the arithmetic mean.

Implementation. *MVSS-Net* is implemented in PyTorch and trained on an NVIDIA Tesla V100 GPU. The input size is 512×512 . The two ResNet-50 used in ESB and NSB are initialized with ImageNet-pretrained counterparts. We use an Adam [15] optimizer with a learning rate periodically decays from 10^{-4} to 10^{-7} . We set the two weights in the combined loss as $\alpha = 0.16$ and $\beta = 0.04$, according to the model performance on a held-out validation set from DEFACTO. We apply regular data augmentation for training, including flipping, blurring, compression and naive manipulations either by cropping and pasting a squared area or using built-in OpenCV inpainting functions [3, 22].

4.2. Ablation Study

For revealing the influence of the individual components, we evaluate the performance of the proposed model in varied setups with the components added progressively.

We depart from FCN-16 without multi-view multi-scale supervision. Recall that we use a DA module for

branch fusion. So for a fair comparison, we adopt FCN-16 with DA, making it essentially an implementation of DANet [10]. The improved FCN-16 scores better than its standard counterpart, *e.g.* UNet [20], DeepLabv3 [5] and DeepLabv3+ [6], see the supplement. This competitive baseline is referred to as *Seg* in Table 3.

Influence of the image classification loss. Comparing *Seg+Clf* and *Seg*, we see a clear increase in specificity and a clear drop in sensitivity, suggesting that adding $loss_{clf}$ makes the model more conservative for reporting manipulation. This change is not only confirmed by lower pixel-level performance, but is also observed in the fourth column of Fig. 6, showing that manipulated areas predicted by *Seg+Clf* are much reduced.

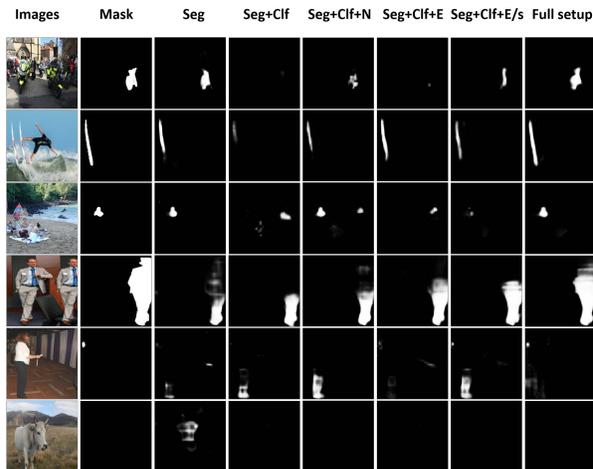


Figure 6. **Pixel-level manipulation detection results of *MVSS-Net* in varied setups.** The test image in the last row is authentic.

Influence of NSB. Since *Seg+Clf+N* is obtained by adding NSB into *Seg+Clf*, its better performance verifies the effectiveness of NSB for improving manipulation detection at both pixel-level and image-level.

Influence of ESB. The better performance of *Seg+Clf+E* against *Seg+Clf* justifies the effectiveness of ESB. *Seg+Clf+E/s* is obtained by removing the Sobel operation from *Seg+Clf+E*, so its performance degeneration in particular on copy-move detection (from 0.405 to 0.382, *cpmv* in Table 3) indicates the necessity of this operation.

ESB versus GSR-Net. *Seg+Clf+G* is obtained by replacing our ESB with the edge branch of GSR-Net. The overall performance of *Seg+Clf+G* is lower than *Seg+Clf+E*. Moreover, there is a larger performance gap on *cpmv* (ESB of 0.405 versus GSR-Net of 0.363). The results clearly demonstrate the superiority of the proposed ESB over the prior art.

Influence of two branch fusion. The full setup, with ESB and NSB fused by dual attention, performs the best,

Setup	Component			Pixel-level manipulation detection (F1)			Image-level manipulation detection				Com-F1	
	loss	ESB	NSB	cmpv.	spli.	inpa.	MEAN	AUC	Sen.	Spe.		F1
Seg	-	-	-	0.453	0.722	0.463	0.546	0.840	0.827	0.620	0.709	0.617
Seg+Clf	+	-	-	0.341	0.673	0.376	0.463	0.858	0.768	0.778	0.773	0.579
Seg+Clf+N	+	-	+	0.393	0.706	0.426	0.508	0.871	0.763	0.821	0.791	0.619
Seg+Clf+E	+	+	-	0.405	0.715	0.435	0.518	0.870	0.773	0.811	0.792	0.626
Seg+Clf+E/s	+	w/o sobel	-	0.382	0.710	0.422	0.505	0.869	0.792	0.789	0.790	0.616
Seg+Clf+G	+	GSR-Net	-	0.363	0.714	0.421	0.499	0.864	0.813	0.779	0.796	0.613
Full setup	+	+	+	0.446	0.714	0.455	0.538	0.886	0.797	0.802	0.799	0.643
Ensemble(N, E)	+	+	+	0.384	0.708	0.437	0.510	0.878	0.731	0.876	0.797	0.622

Table 3. **Ablation study of MVSS-Net.** Training: DEFACTO-84k. Test: DEFACTO-12k. Copy-move, splicing and inpainting are shortened as *cmpv*, *spli* and *inpa*, respectively. Best number per column is shown in **bold**. The top performance of the full setup justifies the necessity of the individual components used in *MVSS-Net*.

showing the complementarity of the individual components. To further justify the necessity of our dual attention based fusion, we make an alternative solution which ensembles *Seg+Clf+N* and *Seg+Clf+E* by model averaging, referred to as *Ensemble(N,E)*. The full setup is better than *Ensemble(N,E)*, showing the advantage of our fusion method¹.

Fig. 6 shows some qualitative results. From the left to right, the results demonstrate how *MVSS-Net* strikes a good balance between sensitivity and specificity. Note that the best pixel-level performance of FCN is due to the fact that the training and test sets are homologous. Next, we evaluate the generalizability of FCN and *MVSS-Net*.

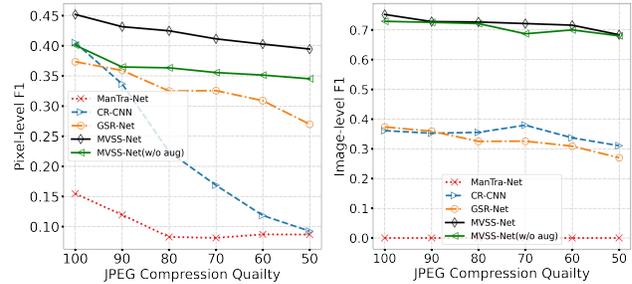
4.3. Comparison with State-of-the-art

Baselines. For a fair and reproducible comparison, we have to be selective, choosing the state-of-the-art that meets one of the following three criteria: 1) pre-trained models released by paper authors, 2) source code publicly available, or 3) following a common evaluation protocol where CASIAv2 is used for training and other public datasets are used for testing. Accordingly, we compile a list of six published baselines as follows:

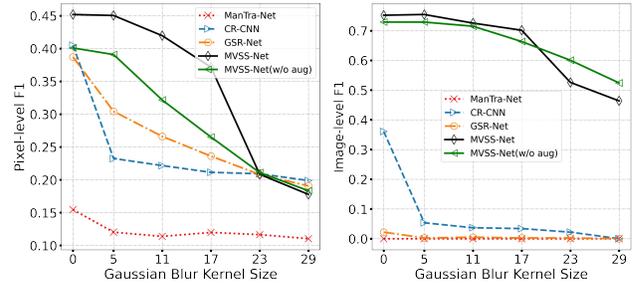
- Models available: HP-FCN [16], trained on a private set of inpainted images², ManTra-Net [26], trained on a private set of millions of manipulated images³, and CR-CNN [27], trained on CASIAv2⁴. We use these models directly.
- Code available: GSR-Net [29], which we train using author-provided code⁵. We cite their results where appropriate and use our re-trained model only when necessary.
- Same evaluation protocol: MFCN [21], RGB-N [30] with numbers quoted from the same team [29].

We re-train FCN (*Seg*) and *MVSS-Net*(full setup) from scratch on CASIAv2.

¹Comparison to fusion by bilinear pooling is in the supplement.
²https://github.com/lihaod/Deep_inpainting_localization
³<https://github.com/ISICV/ManTraNet>
⁴<https://github.com/HuizhouLi/Constrained-R-CNN>
⁵<https://github.com/pengzhou1108/GSRNet>



(a) Performance curves w.r.t. JPEG compression



(b) Performance curves w.r.t. Gaussian Blurs

Figure 7. **Robustness evaluation against JPEG compression and Gaussian Blurs on CASIAv1.**

Pixel-level manipulation detection. The performance of distinct models is given in Table 4. *MVSS-Net* is the best in terms of overall performance. We attribute the clearly better performance of ManTra-Net on DEFACTO-12k to its large-scale training data, which was also originated from MS-COCO as DEFACTO-12k. As *MVSS-Net* is derived from FCN, its superior performance in this cross-dataset setting justifies its better generalizability.

As HP-FCN is specially designed for inpainting detection, we narrow down the comparison to detecting the inpainting subsets in NIST16 and DEFACTO-12k. Again, *MVSS-Net* outperforms HP-FCN: 0.565 versus 0.284 on NIST16 and 0.391 versus 0.106 on DEFACTO-12k.

Method	Optimal threshold per model & testset						Fixed threshold (0.5)					
	NIST	Columbia	CASIAv1	COVER	DEFACTO-12k	MEAN	NIST	Columbia	CASIAv1	COVER	DEFACTO-12k	MEAN
MFCN [21]	0.422	0.612	0.541	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
RGB-N [30]	n.a.	n.a.	0.408	0.379	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
HP-FCN [16]	0.360	0.471	0.214	0.199	0.136	0.276	0.121	0.067	0.154	0.003	0.055	0.080
ManTra-Net [26]	0.455	0.709	0.692	0.772	0.618	0.649	0.000	0.364	0.155	0.286	0.155	0.192
CR-CNN [27]	0.428	0.704	0.662	0.470	0.340	0.521	0.238	0.436	0.405	0.291	0.132	0.300
GSR-Net [29]	0.456	0.622	0.574	0.489	0.379	0.504	0.283	0.613	0.387	0.285	0.051	0.324
FCN	0.507	0.586	0.742	0.573	0.401	0.562	0.167	0.223	0.441	0.199	0.130	0.232
<i>MVSS-Net</i>	0.737	0.703	0.753	0.824	0.572	0.718	0.292	0.638	0.452	0.453	0.137	0.394

Table 4. **Performance of pixel-level manipulation detection.** Best result per test set is shown in bold. All the models are trained on CASIAv2, except for ManTra-Net and HP-FCN.

Method	Columbia				CASIAv1				COVER				DEFACTO-12k			
	AUC	Sen.	Spe.	F1												
ManTra-Net [26]	0.701	1.000	0.000	0.000	0.141	1.000	0.000	0.000	0.491	1.000	0.000	0.000	0.543	1.000	0.000	0.000
CR-CNN [27]	0.783	0.961	0.246	0.392	0.766	0.930	0.224	0.361	0.566	0.967	0.070	0.131	0.567	0.774	0.267	0.397
GSR-Net [29]	0.502	1.000	0.011	0.022	0.502	0.994	0.011	0.022	0.515	1.000	0.000	0.000	0.456	0.914	0.001	0.002
FCN	0.762	0.950	0.322	0.481	0.796	0.717	0.844	0.775	0.541	0.900	0.100	0.180	0.551	0.711	0.338	0.458
<i>MVSS-Net</i>	0.980	0.669	1.000	0.802	0.839	0.615	0.969	0.752	0.731	0.940	0.140	0.244	0.573	0.817	0.268	0.404

Table 5. **Performance of image-level manipulation detection on Columbia, CASIAv1, COVER and DEFACTO-12k.** *Sen.*: sensitivity. *Spe.*: specificity. NIST16, which has no authentic images, is excluded. The default decision threshold of 0.5 is used for all models.

Method	Columbia	CASIAv1	COVER	DEFACTO-12k
ManTra-Net [26]	0.000	0.000	0.000	0.000
CR-CNN [27]	0.413	0.382	0.181	0.198
GSR-Net [29]	0.042	0.042	0.000	0.004
FCN	0.305	0.562	0.189	0.203
<i>MVSS-Net</i>	0.711	0.565	0.317	0.205

Table 6. **Com-F1, the harmonic mean of pixel-level F1 and image-level F1, on four test sets.**

Image-level manipulation detection. Table 5 shows the performance of distinct models, all using the default decision threshold of 0.5. *MVSS-Net* is again the top performer. With its capability of learning from authentic images, *MVSS-Net* obtains higher specificity (and thus lower false alarm rate) on most test sets. Our model also has the best AUC scores, meaning it is better than the baselines on a wide range of operation points.

The overall performance on both pixel-level and image-level manipulation detection is provided in Table 6.

Robustness evaluation. JPEG compression and Gaussian blur are separately applied on CASIAv1. ManTra-Net used a wide range of data augmentations including compression, while CR-CNN and GSR-Net did not use such data augmentation. So for a more fair comparison, we also train *MVSS-Net* with compression and blurring excluded from data augmentation, denoted as *MVSS-Net* (w/o aug). Performance curves in Fig. 7 show better robustness of *MVSS-Net* and *MVSS-Net* (w/o aug).

Efficiency test. We measure the inference efficiency in terms of frames per second (FPS). Tested on NVIDIA Tesla

V100 GPU, CR-CNN, ManTra-Net and GSR-Net run at FPS of 3.1, 2.8 and 31.7, respectively. *MVSS-Net* runs at FPS of 20.1, sufficient for real-time application.

5. Conclusions

Our image manipulation detection experiments on five benchmark sets allow us to draw the following conclusions. For learning semantic-agnostic features, both noise and edge information are helpful, whilst the latter is better when used alone. For exploiting the edge information, our proposed edge-supervised branch (ESB) is more effective than the previously used feature concatenation. ESB steers the network to be more concentrated on tampered regions. Regarding the specificity of manipulation detection, we empirically show that the state-of-the-arts suffer from poor specificity. The inclusion of the image classification loss improves the specificity, yet at the cost of a clear performance drop for pixel-level manipulation detection. Multi-view feature learning has to be used together with multi-scale supervision. The resultant *MVSS-Net* is a new state-of-the-art for image manipulation detection.

Acknowledgements. This research was supported by NSFC (U1703261), BJNSF (4202033), the Fundamental Research Funds for the Central Universities and the Research Funds of Renmin University of China (No. 18XNLG19), and Public Computing Cloud, Renmin University of China. This work was initially inspired by the Security AI Challenge: Forgery Detection on Certificate Image, Alibaba Security.

References

- [1] J. Bappy, A. Roy-Chowdhury, J. Bunk, L. Nataraj, and B. Manjunath. Exploiting spatial structure for localizing manipulated image regions. In *ICCV*, 2017. 1, 3
- [2] B. Bayar and M. Stamm. Constrained convolutional neural networks: A new approach towards general purpose image manipulation detection. *IEEE Transactions on Information Forensics and Security*, 13(11):2691–2706, 2018. 2, 3, 4
- [3] M. Bertalmio, A. Bertozzi, and G. Sapiro. Navier-stokes, fluid dynamics, and image and video inpainting. In *CPVR*, 2001. 6
- [4] L. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. Yuille. DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):834–848, 2018. 1
- [5] L. Chen, G. Papandreou, F. Schroff, and H. Adam. Rethinking atrous convolution for semantic image segmentation. In *CVPR*, 2017. 6
- [6] L. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. 6
- [7] J. Dong, W. Wang, and T. Tan. CASIA image tampering detection evaluation database. <http://forensics.idealtest.org>, 2010. 1, 2, 5, 6
- [8] J. Dong, W. Wang, and T. Tan. CASIA image tampering detection evaluation database. In *ChinaSIP*, 2013. 1, 2, 5, 6
- [9] J. Fridrich and J. Kodovsky. Rich models for steganalysis of digital images. *IEEE Transactions on Information Forensics and Security*, 7(3):868–882, 2012. 2, 3
- [10] J. Fu, J. Liu, H. Tian, Y. Li, Y. Bao, Z. Fang, and H. Lu. Dual attention network for scene segmentation. In *CVPR*, 2019. 5, 6
- [11] O. Gafni and L. Wolf. Wish you were here: Context-aware human generation. In *CVPR*, 2020. 1
- [12] H. Guan, M. Kozak, E. Robertson, Y. Lee, A. N. Yates, A. Delgado, D. Zhou, T. Kheyrkhan, J. Smith, and J. Fiscus. MFC datasets: Large-scale benchmark datasets for media forensic challenge evaluation. In *WACV Workshop*, 2019. 2, 5, 6
- [13] J. Hsu. Columbia uncompressed image splicing detection evaluation dataset. <https://www.ee.columbia.edu/ln/dvmm/downloads/AuthSplicedDataSet/AuthSplicedDataSet.htm>, 2009. 2, 5, 6
- [14] X. Hu, Z. Zhang, Z. Jiang, S. Chaudhuri, Z. Yang, and R. Nevatia. SPAN: Spatial pyramid attention network for image manipulation localization. In *ECCV*, 2020. 1, 2, 3
- [15] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *Computer Science*, 2014. 6
- [16] H. Li and J. Huang. Localization of deep inpainting using high-pass fully convolutional network. In *ICCV*, 2019. 2, 3, 5, 7, 8
- [17] T. Lin, M. Maire, S. Belongie, J. Hays, and C. Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014. 5
- [18] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(4):640–651, 2015. 1
- [19] G. Mahfoudi, B. Tajini, F. Retraint, F. Morain-Nicolier, and M. Pic. DEFACTO: Image and face manipulation dataset. In *EUSIPCO*, 2019. 1, 5, 6
- [20] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 6
- [21] R. Salloum, Y. Ren, and C. Kuo. Image splicing localization using a multi-task fully convolutional network (MFCN). *Journal of Visual Communication and Image Representation*, 51(feb.):201–209, 2017. 1, 2, 3, 5, 6, 7, 8
- [22] A. Telea. An image inpainting technique based on the fast marching method. *Journal of Graphics Tools*, 9(1):23–34, 2004. 6
- [23] L. Verdoliva. Media forensics and deepfakes: An overview. *IEEE Journal of Selected Topics in Signal Processing*, 14(5):910–932, 2020. 1
- [24] Q. Wei, X. Li, W. Yu, X. Zhang, Y. Zhang, B. Hu, B. Mo, D. Gong, N. Chen, D. Ding, and Y. Chen. Learn to segment retinal lesions and beyond. In *ICPR*, 2020. 2, 5
- [25] B. Wen, Y. Zhu, R. Subramanian, T. T. Ng, and S. Winkler. COVERAGE – A novel database for copy-move forgery detection. In *ICIP*, 2016. 1, 5, 6
- [26] Y. Wu, W. AbdAlmageed, and P. Natarajan. ManTra-Net: Manipulation tracing network for detection and localization of image forgeries with anomalous features. In *CVPR*, 2019. 1, 2, 3, 5, 7, 8
- [27] C. Yang, H. Li, F. Lin, B. Jiang, and H. Zhao. Constrained R-CNN: A general image manipulation detection model. In *ICME*, 2020. 1, 2, 3, 5, 7, 8
- [28] C. Yu, J. Wang, C. Peng, C. Gao, G. Yu, and N. Sang. Learning a discriminative feature network for semantic segmentation. In *CVPR*, 2018. 2
- [29] P. Zhou, B. Chen, X. Han, M. Najibi, and L. Davis. Generate, segment, and refine: Towards generic manipulation segmentation. In *AAAI*, 2020. 1, 2, 3, 4, 5, 6, 7, 8
- [30] P. Zhou, X. Han, V. I. Morariu, and L. S. Davis. Learning rich features for image manipulation detection. In *CVPR*, 2018. 2, 3, 5, 6, 7, 8