

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

GCA-Net : Utilizing Gated Context Attention for Improving **Image Forgery Localization and Detection**

Sowmen Das^{1†} Md. Saiful Islam^{2¶} Md. Ruhul Amin^{3§}

[†]Shahjalal University of Science and Technology, Bangladesh [¶]University of Alberta, Canada [§]Fordham University, USA 1 sowmendipta@gmail.com 2 mdsaifu1@ualberta.ca ³mamin17@fordham.edu

Abstract

Forensic analysis of manipulated pixels requires the identification of various hidden and subtle features from images. Conventional image recognition models generally fail at this task because they are biased and more attentive towards the dominant local and spatial features. In this paper, we propose a novel Gated Context Attention Network (GCA-Net) that utilizes non-local attention in conjunction with a gating mechanism in order to capture the finer image discrepancies and better identify forged regions. The proposed framework uses high dimensional embeddings to filter and aggregate the relevant context from coarse feature maps at various stages of the decoding process. This improves the network's understanding of global differences and reduces false-positive localizations. Our evaluation on standard image forensic benchmarks shows that GCA-Net can both compete against and improve over state-of-the-art networks by an average of 4.7% AUC. Additional ablation studies also demonstrate the method's robustness against attributions and resilience to false-positive predictions.

1. Introduction

Image manipulation is the act of altering an image's content using different editing techniques. Examples of such manipulation include, content removal [24], face-swapping [22], attribute changing [39], *etc.* There are three primary types of manual image forgery techniques [6] – splicing, copy-move, and inpainting. The majority of these manipulation scenarios consist of copying and pasting some pixels from a source image onto a target image, as shown in. Fig. 1. Forensic algorithms play the important role of determining whether an image is authentic or not. Further research in developing better detection mechanisms is critical in today's age of social media, as any fake news can spread rapidly and be used to foment panic and propaganda.

Numerous methods have been proposed over the years for image forgery localization and detection (IFLD). In recent years the use of deep CNN architectures capable of



Prediction

Figure 1. The first and second columns contain examples of authentic and manipulated images. The third column shows the output generated by our proposed network, highlighting the manipulated regions of the tampered images.

learning intrinsic features has gained popularity [2]. However, unlike traditional image classification, forgery detection involves identifying hidden manipulation traces from the image rather than the apparent spatial content. Existing works that utilize a variety of sequential frameworks [17, 40, 54] suffer from feature attenuation due to the subtlety of the forgery artifacts. Although some methods try to utilize feature hierarchies through recursive pooling [1,21], they are mostly bound to local neighborhoods.

Majority of the existing manipulation localization networks [13,21,54] perform a pixel-level binary classification of each individual pixel in order to generate the localization map. Image-level classification is consequently done depending on the percentage of identified pixels. This simplification results in two notable drawbacks -i) Since these methods view each pixel as a single data point, they fail to utilize important region characteristics such as noisy boundaries and object artifacts that can help in discerning region differences, ii) Despite authentic images not containing manipulation traces, methods that were trained using this process exhibit erroneous predictions when trying to localize untampered images; shown in Fig. 2.

In order to address the shortcomings of existing networks, we propose the Gated Context Attention Network



Untampered Image

False Predictions

Figure 2. Examples of false predictions from different networks. The first column contains authentic images and the next two contain results of the following networks - (a) Bappy *et al.* [1], (b) Zhou *et al.* [58], (c,e) Wu *et al.* [54], (d) Wu *et al.* [52], (f) Wu *et al.* [53]. We can see that the networks highlight a lot of erroneous pixels for untampered images.

(GCA-Net), which is composed of an improved feature encoder and a densely connected attention decoder. Image manipulation alters not only the visible content of an image but also the underlying signal and noise characteristics. So, in order to utilize these multi-modal features, we combine EfficientNet [44] with an improved Error Level Analysis (ELA) [49] module and other steganalysis layers to identify the compression and noise artifacts. However, since every image contains some intrinsic noise patterns [13], in order to localize the manipulations, we must correlate the local noise difference to the global image fingerprint. Although previous works [21, 54, 59] have attempted to use these signal features, their inattention to the global similarities resulted in generating noisy false positives.

Our proposed GCA-Net uses a novel gated context attention which performs two distinct computations. Firstly, it aggregates multiple cross-layer feature embeddings and applies a global context attention to identify differences between local and global representations. This improves the network's understanding of long-range dependencies between the various pixel regions. Secondly, the low-level context features are filtered using an attention gate to eliminate redundant spatial and object data. By propagating only the necessary forensic features, we are able to minimize attenuation and improve the identification of boundary artifacts. Furthermore, the dense network structure of Unet++ [60] facilitates information exchange between intermediate layers, thus reducing feature loss and improving convergence. This design also increases the network's robustness to post-processing attributions. Our evaluation on the popular CASIA [15], NIST [19], and IMD [33] benchmarks show performances comparable to existing SOTA methods under constrained resources. GCA-Net outperforms several existing methods by 4% - 6%, successively generating better localization and fewer false positives. To summarize the contributions of our work:

- We propose GCA-Net consisting of a novel Gated Context Attention module that enables efficient modelling of long-range dependencies and improves global context representation necessary for manipulation localization.
- We substantiate the impact of global context and nonlocality in detecting image forgeries, facilitating future research in this field.
- We illustrate the effects of training strategies for reducing false positives and improving image-level predictions.
- We show that our method outperforms the state-of-the-art performance on standard datasets for image manipulation detection and localization.

2. Related Works

Depending on the type of forgery, different clues and artifacts can be found in an image to determine whether it is authentic or not. These include, compression noise [47], PRNU sensor information [11, 12], camera model information [7, 35], local noise features [13, 42], etc. Earlier forensic methods used algorithms to detect these noise and signal properties using carefully designed filters. Among them, the SRM filters used in [59] have shown success in identifying hidden noise patterns from images.

A number of neural networks have also been proposed for both independent forgery detection [26, 31, 36, 45], and localization [7, 32, 53]. Object detection networks using RCNN, region proposal modules, and bounding box identification have shown to be effective for manipulation localization [4, 50, 59]. Consequently, segmentation networks like Unet and DeepLab have also been used [5,9]. Furthermore, the use of constrained layers have also shown effectiveness in improving input representation and minimizing the effects of dominant spatial features [3, 56]. ManTra-Net [54] uses a simple VGG [41] network with a Z-pooling method to localize anomalous features. This work was extended by SPAN [21] to further model the spatial correlation via local self-attention blocks and hierarchical pyramid propagation. However, both networks fail to utilize the correlation of global context and multi-scale features. In a recent work [27], PSCC-Net was proposed that tries to address these problems using channel correlation. They showed that multi-scales attention could be leveraged to improve manipulation detection.

While self-attention was originally introduced for language modeling [46], this seminal work has been proven to improve long-range feature representation across a va-



Figure 3. The architecture of GCA-Net. The input image first passes through feature suppression layers similar to [21,54] that extracts multi-modal compression and noise level features. These are then propagated through the EfficientNet encoder backbone consisting of five stages $X^{0,0} \to X^{4,0}$. Each node of the dense decoder $X^{i,j}$ contains the Gated Context Attention block that takes in the i^{th} layer features $F_l = [X^{i,0}, ..., X^{i,j-1}]$ shown using blue solid and dashed arrows, and the gating feature F_g from $(i + 1)^{th}$ layer shown using orange arrows. Due to the bottom-up architecture of the decoder, the outputs from zeroth layer $X^{0,1\to4}$ are taken and using deep supervision [60] the final localization map is generated. We use separate classification and localization losses for multi-task learning and optimization.

riety of domains. Several works have utilized variations of attention for forgery detection [10, 14, 30]. Vision tasks primarily use two types of attention mechanisms: channelattention [20] and spatial-attention [48]. Spatial attention using non-local blocks, also termed dot-product attention, can be considered a generic framework for non-linear correlation. However, this method suffers from quadratic complexities as the attention maps are computed for every pixel pair. An improved context attention was introduced in [8] that replaces the quadratic operation by simplifying the context modeling framework. In our work, we use this global context framework in conjunction with a gated attention mechanism to explore discrepancies in spatial and feature channels concurrently. Attention gating has previously been used for various tasks including graph learning [57], language processing [55], medical image segmentation [34], etc. We introduce gating for forgery detection in order to reinforce and filter coarse-level features during the upscaling and decoding phases. This leads to improvements in information sharing by only propagating the necessary forensic signals through the network.

3. Proposed Method

3.1. Overview

As illustrated in Fig. 3, GCA-Net is a multi-branch dense encoder-decoder network. The model is comprised of three parts: a feature encoder, a dense decoder, and two heads for classification and localization. An input image first passes through a series of content suppression layers that generate multi-domain semantic features. While standalone CNN layers are capable of extracting features from images, they are localized and biased towards the dominant spatial features rather than the manipulation traces. So, in addition to the constrained layers in [54], we use an Error Level Analysis (ELA) module to extract signal noise and compression artifacts. The encoded backbone features and intermediate layer outputs are used by the attention decoder to produce the final localization map. Each decoder block is composed of a GCA layer followed by a series of convolution layers for feature accumulation and upscaling. Additionally, the encoder features also pass through a fully connected layer to generate an image-level probability score. The simultaneous classification and localization utilizes a multi-task learning approach for better feature representation. This plays a significant role in reducing false positives and improving generalizability. Additional details regarding the implementation of ELA, and the effect of each layer on the network's performance can be found in Suppl. A1.

3.2. Dense Feature Decoder

For the decoder branch, we use the dense Unet++ [60] architecture that utilizes a multi-scale bottom-up architecture for feature propagation. We denote the upper layers of the decoder as *coarse* layers, and the lower layers as *fine* based on the selectivity of their feature maps. The skip connections between the layers allow the flow of multi-modal semantic features, thus improving gradient flow and min-

imizing feature loss. Moreover, since the decoder samples the intermediate layers at multiple scales, it improves global context representations.

Features from an input $I \in \mathbb{R}^{3 \times H \times W}$ are sampled at five intermediate scales $H/k \times W/k$, where $k \in$ $\{2, 4, 8, 16, 32\}$. Each gray circle in Fig. 3 represents a decoder node denoted $X^{i,j}$, where *i* indexes the sequential down-sampling layers of the encoder and *j* indexes a decoder block along the skip pathway at the *i*th layer. The output $y^{i,j}$ at any node is computed as,

$$y^{i,j} = \mathcal{C}\left(\vartheta\right) \tag{1}$$

$$\vartheta = \Theta\left(\Omega\left(\left[y^{i,k}\right]_{k=0}^{j-1}\right), \mathcal{U}\left(y^{i+1,j-1}\right)\right)$$
(2)

where $C(\cdot)$ is a series of convolution operations followed by a ReLU activation, $\Theta(\cdot)$ is the GCA layer, $\Omega(\cdot)$ is a nonlocal context block, $U(\cdot)$ denotes an up-sampling layer, and $[\cdot]$ denotes the concatenation layer.

The bottom-up architecture ensures that at each node features from all lower scales are accumulated. For instance, $X^{3,1}$ is reinforced by $X^{4,0}$ and $X^{3,0}$. Similarly, computation of the final node $X^{0,4}$ aggregates features from all previous layers and nodes $X^{4,0}$ up to $X^{0,3}$. Thus, each node can determine which features are most relevant reinforced by the finer lower-level features and propagate it to upper branches. In contrast to standard encoder-decoder networks, this multi-scale fusion enables easier identification of the essential attributes without relying on a single previous output.

3.3. Gated Context Attention (GCA)

Attention mechanisms are used to modulate learned features according to their relative significance. The GCA operation shown in Fig. 4 is divided into two primary stages: 1) Global Context Pooling and 2) Attention Gating.

Deep convolution stacks tend to obfuscate global pixelto-pixel relationships due to their nature of locality [48]. Non-local blocks solve this problem using attention weights and aggregating information from other points to reinforce the features of a query position. Global context modeling [8] is an improved attention framework for identifying long range dependencies between feature maps. In the first stage, we compute the global context from the concatenated features of the current level. For instance, to compute the context of node $X^{0,3}$, we aggregate $[X^{0,0}, X^{0,1}, X^{0,2}]$ which is denoted as F_l or *layer features*. These are coarse feature maps containing a greater amount of global information than the subsequent finer layers. Because identification of manipulation features is based on detecting changes between a group of pixels and their surroundings, using these global contexts allows the model to recognize the differences between altered regions. We can rewrite Eq. (2) as, $\vartheta = \Theta(\Omega(F_l), F_g)$, where, $F_l = [y^{i,k}]_{k=0}^{j-1}$ is the concate-



Figure 4. The structure of Gated Context Attention block. The blue line shows the flow concatenated features F_l of i^{th} decoder level and the orange line denotes the upsampled gating feature F_g from $(i + 1)^{th}$ level as referred in Fig. 3

nation of i^{th} level features, and $F_g = \mathcal{U}(y^{i+1,j-1})$ is the up-sampled feature of $(i+1)^{th}$ level.

The context block Ω has three steps - i) attention pooling, ii) feature transform, and iii) feature fusion. For the pooling step we take $F_l \in \mathbb{R}^{C_l \times H \times W}$ and pass it through a 1×1 convolution to reduce the channels to $C_l \times 1 \times 1$. It groups the features of all positions via weighted averaging to obtain the global context vector. This is similar to the Global Average Pooling of Squeeze-Excitation (SE) layer [20]. The pooled vector is then passed through a bottleneck block with a factor r to capture the channel-wise dependencies. This is the transform step. The reduction and expansion is analogous to the excitation operation of SE block. Finally, the fusion step aggregates the context features to the features of each input position using a broadcast element-wise addition.

The second stage of GCA is the attention gating which is used to filter the coarse level features. Attention Gates (AG) identify salient image regions and prune feature responses to retain only relevant activations. As a result, signals

from irrelevant background regions and spatial contexts are gradually suppressed. AGs generate a coefficient matrix, $A_a \in [0,1]$, which is multiplied with the input to alter the scale of the activations. Each AG learns to focus on a subset of target structures reinforced by the downstream features. The coarse maps of the encoder represent global relationships, while the downstream layers identify finer discriminating features. Gating uses these finer embeddings to disambiguate irrelevant and noisy responses within the coarse features. At each decoder node $X^{i,j}$, the gating feature F_g is the upsampled feature propagated by the node $X^{i+1,j-1}$ highlighted in orange. In order to obtain the gating coefficient via additive attention, we first transform the context feature $\Omega(F_l) \in \mathbb{R}^{C_l \times H \times W}$ and gating feature $F_g \in \mathbb{R}^{C_g \times H \times W}$ to an intermediate vector $F_{int} \in \mathbb{R}^{C_{int} \times H \times W}$. A non-linear transformation layer $W_{\zeta} = ReLU(Conv(\cdot))$ is then used to resample this vector. This produces the final attention matrix $A_g = W_{\zeta}(\Omega(F_l) \oplus F_G) \in \mathbb{R}^{1 \times H \times W}$. Finally, the attention matrix is multiplied with the input feature stack F_l to generate $\vartheta \in \mathbb{R}^{C_l \times H \times W}$, the GCA output. The resulting gated features are sent through a series of convolution and activation layers $\mathcal{C}(\cdot)$ that perform the decoding operation at the particular node. Additional spatial and channel attention layers [37] can be added after this stage to further modulate the learnt features.

4. Experiments

4.1. Datasets

We follow the evaluation protocols in [54, 59] for training and validation. We train our model on four types of data - splicing, copy-move, inpainting, and authentic images; collected from Dresden [18], MS COCO synthetic [1], Defacto [29], and IMD-Real [33] datsets. For both pretrained and finetuned evaluation we use the four standard datasets: CASIAv2 [15], NIST16 [19], COVERAGE [51] and IMD-2020 [33], following the training/testing split described in [27]. In total, we used ~ 170 k images for training, with roughly equal class distribution. This amount is significantly lower compared to existing SOTA methods like SPAN [21], MantraNet [54], and PSCC-Net [27] that use upwards of 500k~1M images. We did not use a larger dataset due to resource and accessibility constraints. The majority of our computation was conducted on an Nvidia 1080 Ti GPU. However, despite the smaller train set, our model outperforms these methods in multiple experiments.

4.2. Loss Function

The train data consists of input images $I \in \mathbb{R}^{3 \times H \times W}$ and binary ground-truth masks $M \in [0, 1]^{1 \times H \times W}$. GCA-Net was trained using a multi-task loss function combining detection and localization losses. Most traditional methods [21, 27, 54] train the localization network using Binary Cross-Entropy (BCE) loss (eq. 3). BCE works well for classification problems with lots of data and a balanced dataset because it weighs all predictions equally.

$$\mathcal{L}_{BCE} = -\frac{1}{N} \sum_{i=0}^{N} y_i \cdot \log\left(\hat{y}_i\right) \tag{3}$$

The loss function minimizes the distance between the predicted and ground truth value for all predictions. For example, if a fake pixel is predicted with a probability of 0.7, BCE will try to bring it closer to 1 in order to minimize the loss. Thus, for unbalanced datasets, predictions become skewed toward a particular class. Since we are working with a small and unbalanced dataset, we opt to use a combination of Dice [43], and Focal [25] loss. We compute the dice loss using the following equation,

$$\mathcal{L}_{DSC}(P,G) = -\log\left(\frac{2 \cdot |P \cap G| + \epsilon}{|P| + |G| + \epsilon}\right) \tag{4}$$

For a prediction mask P and ground-truth G both having dimensions $1 \times H \times W$, the numerator calculates the intersection of the regions, and the denominator measures the union. Dice loss minimizes the distance between the predicted and true regions. Rather than calculating loss for each individual pixel, it improves prediction for the entire region. This is critical in lowering our false positive rate. Since BCE tries to minimize loss for each pixel independently, the network always produces incorrect predictions for a certain number of pixels.

Even though Dice loss improves false positive and region overlaps, it falls short for small forged regions. This is because as the regions become very small, |G| becomes small. So if the model does not predict anything at all, *i.e.*, if $|P| \rightarrow 0$, then $|P \cap G| \rightarrow 0$, and the total loss decreases. Thus, dice loss alone is ineffective in these instances. To overcome this problem, we combine Focal loss with Dice loss. Focal loss is an improvement over BCE loss with an additional temperature parameter to account for overconfident predictions. Focal loss is calculated as,

$$\mathcal{L}_{FL}(P) = -\alpha_{t} \left(1 - P\right)^{\gamma} \log\left(P\right) \tag{5}$$

The temperature γ controls the weighting of each prediction. When $\gamma > 0$ is used, weak predictions are weighted more heavily. So, the network focuses on improving a prediction of 0.3 towards 1 rather than trying to improve a prediction of 0.7. Since we are trying to correctly classify fake pixels, a larger prediction P results in a smaller $(1 - P)^{\gamma}$, thereby reducing the overall loss. So, the network will try to classify weaker predictions more accurately. Combining all of the above, our final loss function becomes,

$$\mathcal{L} = w_c \cdot \mathcal{L}_{CLS} + w_d \cdot \mathcal{L}_{DSC} + w_f \cdot \mathcal{L}_{FL} \tag{6}$$

Loss Function	AUC	F1
\mathcal{L}_{BCE}	78.33	60.29
\mathcal{L}_{DSC}	84.12	64.41
$\mathcal{L}_{DSC} + \mathcal{L}_{FL}$	86.05	65.50
$\mathcal{L}_{DSC} + \mathcal{L}_{FL} - Reduction$ [38]	86.81	67.25
$\mathcal{L}_{CLS} + \mathcal{L}_{DSC} + \mathcal{L}_{FL} - Reduction$	88.95	72.39

 Table 1. Localization performance of GCA-Net under different loss functions on CASIAv2 validation set.

 $\mathcal{L}_{CLS} = \mathcal{L}_{BCE}$ is the classification loss for the encoder's prediction. By combining classification loss, the network can optimize both the encoder and decoder at the same time. Each loss parameter can be adjusted independently. We use $\epsilon = 10^{-7}, \gamma = 2, w_c = 1, w_d = 1.10$, and $w_f = 1.15$ to train our final network. In Table 1 we report the results of different configuration of loss functions on the CASIAv2 validation set. To quantify localization performance, we use pixel-level AUC following previous works [27, 54, 59] and Dice score *i.e.*, pixel F1 score.

4.3. Ablation Study

We follow the ablation study in [8] to compare changes of parameters and configuration of the network. The following experiments are done on the CASIAv2 validation set, and we report the pixel-level AUC and F1 score.

Block Design: For the choice of long range context modeling we compare other existing frameworks - Non-local (NL) block, Simplified Non-local (SNL) block, SE block, and the Global Context block, placed before every decoder node. Table 2(a) shows all context frameworks achieve better performance over baseline. NL and SNL blocks are quite similar, while GCA blocks with comparably fewer parameters yield the best performance.

Bottleneck Design: The effects of each component in the bottleneck section are shown in Table 2(b). w/o ratio uses a single 1×1 convolution as a transform which has higher parameters and achieves the best performance. Even though r16+ReLU has fewer parameters they are harder to optimize. Thus Layer Norm (LN) is used to ease optimization, leading to performance similar to w/o ratio.

Bottleneck Ratio: The bottleneck is used to reduce redundancy in parameters and provide a trade-off between parameter and performance. The bottleneck ratio r controls the amount of feature compression. Table 2(c) shows that the network's performance improves consistently as the ratio decreases. We use a bottleneck ratio r = 4 which has a good balance of performance and parameters.

Block Positions: We determine whether the placement of the GCA blocks have an effect on performance. Various placement positions are illustrated in Suppl. A2. As seen in Table 2(e), this effects performance only to a small degree.

(a) Block Design			
	AUC(%)	F1	
baseline (without GCA)	85.2	68.4	
non-local	86.1	69.8	
simplified non-local	86.9	70.5	
global context	87.4	71.7	
GCA	88.9	72.4	
(b) Bottleneck	Design		
baseline (without GCA)	85.2	68.4	
w/o ratio	88.9	72.5	
r16 (ratio 16)	87.3	70.4	
r16+ReLU	87.3	70.4	
r16+LN+ReLU	88.3	72.6	
(c) Bottleneck Ratio			
baseline (without GCA)	85.2	68.4	
ratio 4	88.8	72.8	
ratio 8	88.5	72.8	
ratio 16	88.3	72.6	
ratio 32	87.9	71.7	
(e) Block Positions			
baseline (without GCA)	85.2	68.4	
all decoder	88.3	72.6	
only end nodes	87.1	70.5	
only intermediates	87.8	71.9	
only top nodes	86.3	69.4	

Table 2. Ablation study of the GCA block for pixel-level localization on the CASIAv2 validation set.

The best results are obtained by placing the attention block before each node.

Additional ablation experiments regarding backbone choice, and training specifics are provided in Suppl. A3.

5. Comparison and Evaluation

We compare GCA-Net against existing SOTA architectures for both pre-trained and finetuned evaluation. The pre-trained model was selected based on the best validation score on the train set. We report existing values as mentioned in [27]. For finetuned evaluation we use unseen test splits generated following the evaluation process in [59].

In Tables 3, 4 we compare GCA-Net to existing methods. Both the pre-trained and fine-tuned comparisons demonstrate that GCA-Net outperforms all other methods on the CASIAv2 and IMD-2020 datasets and is comparable on the NIST and COVERAGE datasets. On CASIA and IMD, we see an improvement of 5.4% and 4.46% respectively over PSCC-Net. GCA-Net achieves an AUC of 95.3 on NIST-16, trailing behind PSCC-Net by only 4%. However, we surpass every network on NIST-16 in terms of F1 score. This is because of our fine-tuning the loss functions. Pixellevel F1 score measures the region overlap of the prediction and ground-truth. Since we optimized the network using

Method	CASIAv2	COVERAGE	NIST16	IMD2020
MantraNet [54]	81.7	81.9	79.5	74.8
SPAN [21]	79.7	92.2	84.0	75.0
PSCC-Net [27]	82.9	84.7	85.5	80.6
GCA-Net	87.1	83.1	85.2	81.3

Table 3. Comparison of localization AUC against existing methods using their pre-trained models.

CASIAv2	COVERAGE	NIST16	IMD2020
61.3 / 21.4	58.3 / 22.2	42.9 / 23.6	-
61.2 / 26.3	58.7 / 26.9	48.7 / 28.5	58.6/-
52.2 / 20.7	48.5 / 19.0	50.1 / 17.4	48.7 / -
79.5 / 40.8	81.7 / 43.7	93.7 / 72.2	-
83.8 / 38.2	93.7 / 55.8	83.6 / 29.0	75.0/-
87.5 / 55.4	94.1 / 72.3	99.6 / 81.9	80.6 / -
92.2 / 71.2	87.4 / 69.5	95.3 / 84.5	86.4 / 42.6
	CASIAv2 61.3 / 21.4 61.2 / 26.3 52.2 / 20.7 79.5 / 40.8 83.8 / 38.2 87.5 / 55.4 92.2 / 71.2	CASIAv2 COVERAGE 61.3 / 21.4 58.3 / 22.2 61.2 / 26.3 58.7 / 26.9 52.2 / 20.7 48.5 / 19.0 79.5 / 40.8 81.7 / 43.7 83.8 / 38.2 93.7 / 55.8 87.5 / 55.4 94.1 / 72.3 92.2 / 71.2 87.4 / 69.5	CASIAv2 COVERAGE NIST16 61.3 / 21.4 58.3 / 22.2 42.9 / 23.6 61.2 / 26.3 58.7 / 26.9 48.7 / 28.5 52.2 / 20.7 48.5 / 19.0 50.1 / 17.4 79.5 / 40.8 81.7 / 43.7 93.7 / 72.2 83.8 / 38.2 93.7 / 55.8 83.6 / 29.0 87.5 / 55.4 94.1 / 72.3 99.6 / 81.9 92.2 / 71.2 87.4 / 69.5 95.3 / 84.5

Table 4. Evaluation against existing fine-tuned models for pixellevel localization AUC/F1 score on unseen test splits.

Dice loss, its region identification is better than the existing models. We rank third in COVERAGE, behind SPAN and PSCC-Net. COVERAGE includes samples with very small shifts of copied regions followed by contrast correction and edge blurring. Our train data, which is composed entirely of publicly available datasets, is free of such perturbations, resulting in a difference in the train-test distribution. This limitation can be overcome by training on synthetic copymove data supplemented with adversarial examples. Qualitative examples of localization are shown in Fig. 5. Additional examples comparing localization for both authentic and forged images against MantraNet is provided in the supplementary.

5.1. Detection Performance

To analyze the image-level detection performance, we compare pretrained GCA-Net to SOTA architectures using the metrics reported in [27]. We use a detection dataset containing 511 forged and 749 real images taken from CASIAv2. As can be seen from Table 5, GCA-Net significantly outperforms all other models. This is because it was equipped with a dedicated classification head that was trained particularly to recognize images with forged content. In comparison, existing approaches identify the image class by counting the number of manipulated pixels discovered in the localization result.

Method	Image-Level F1 Score
MantraNet	56.69
SPAN	63.48
PSCC-Net	66.88
GCA-Net (pretrained)	85.51

Table 5. Comparison of image-level detection performance onCASIAv2 detection set against other methods.

5.2. False Positive Evaluation

One of the primary contributions of our proposed method is the reduction of false positives in localization. In order to evaluate the degree of improvement against other networks, we calculate FPR or False Positive Rate for both authentic and manipulated images. FPR quantifies the proportion of data in a sample that is incorrectly identified. In our case, a false positive prediction occurs when a pixel that should be authentic, i.e., 0, is classified as fake, i.e., 1. FPR measures the fraction of incorrect predictions made against the entire set of pixels in an image $FPR = \frac{False Positive Pixels}{Total Number of Pixels}$. However, because the number of incorrect predictions is relatively small in comparison to the total number of pixels in an image, this value can become extremely small and difficult to interpret. Thus, we use the -log(FPR) as our evaluation metric. Increased values indicate better performance and lower false positives.

		ManTraNet	GCA-Net
CASIA	Authentic (A)	5.39	9.88
	Tampered (T)	4.77	5.09
	Combined (A+T)	4.95	5.97
IMD	Authentic (A)	4.57	6.10
	Tampered (T)	4.43	5.65
	Combined (A+T)	4.66	6.93
Authenti	c (CASIA + IMD)	4.94	8.83
Tampere	d (CASIA + IMD)	4.56	5.30

Table 6. Comparison of false positive, -log(FPR), for authentic and manipulated images from CASIAv2 and IMD2020 test sets. Higher values denote less false positives. The experiments were done separately on authentic and tampered images, as well as a combination of both from the two datasets.

To conduct the evaluation, we created a test set using 200 authentic and 100 tampered images taken independently from both CASIAv2 and IMD2020 datasets. We used the pretrained model for GCA-Net, and the publicly available implementation for ManTraNet. As illustrated in Table 6, GCA-Net consistently shows better performance. For authentic images in both CASIA and IMD, GCA-Net's score is almost double that of the other methods. Moreover, due to the logarithmic nature of the metric, fairly small differences denote significant real world improvement. Additional qualitative samples of authentic predictions are provided in the supplementary.

5.3. Robustness Analysis

We examine the performance of our proposed method against various attacks/post-processing to further verify its efficacy and robustness. For this purpose, we degrade images from the NIST16 test set using the distortion settings in [27]. These include Gaussian Blur with kernel size k,



Figure 5. Qualitative examples of manipulation localization with their respective groundtruths from multiple datasets.



Figure 6. Comparison of model robustness against different postprocessing methods.

JPEG Compression with a quality factor q, and Additive Gaussian Noise using standard deviation σ . The reported metrics are calculated using the pretrained GCA-Net. In Fig. 6, we can see that our model outperforms existing approaches against various post-processing attacks.

5.4. Limitations

In our experiments we found that GCA-Net might fail when the manipulation region is very large compared to authentic pixels. Fig. 7(a) contains a sample from NIST-16 where the entire white region is manipulated. Although GCA-Net could detect the discrepancy between the authentic and forged regions, it is not confident about the prediction. It highlighted the centre portion denoting that the region is different from the surrounding pixels. For a second example, we tested our network similar to [54] to check for manual assistance applicability. The initial image in Fig. 7(b) was of dimension 1024×1520 having a small forged region. The network failed to locate the region when tested with the entire image. After that, we cropped the image around the forged location and again tested the cropped image. This time the network was able to identify the manipulated region. This indicates that GCA-Net could be used as a computer-aided tool.



Figure 7. (a) Example of a failed prediction. The groundtruth (2nd image) shows white pixels as forged and black as authentic. The prediction (3rd image) identified discrepancies between the two regions, but inverted the classes. (b) The left pair shows example of a failed localization for a large resolution image. In the right pair, for the same image when tested with a cropped portion, the network successfully located the forged pixels.

6. Conclusion

In this paper, we introduced a novel Gated Context Attention Network (GCA-Net) for detecting and localizing image forgeries. Our proposed network uses a gated attention block to utilize the global context features together with the region attributes to localize manipulated pixels. The proposed attention framework improves longrange dependency modeling and reduces attenuation of the hidden forensic features. This paper illustrated the problems surrounding existing methods and how they might be addressed with better feature representation and training strategies. As demonstrated by our results, GCA-Net outperforms existing SOTA architectures on multiple benchmark datasets by upto 6% with significantly lower false positives. In the future, we will further improve the method in handling large resolution images and work on reducing the limitations. We will also explore the model's viability in detecting deep learning-based deepfake forgeries and other semantic segmentation tasks.

References

- Jawadul H Bappy, Cody Simons, Lakshmanan Nataraj, BS Manjunath, and Amit K Roy-Chowdhury. Hybrid lstm and encoder-decoder architecture for detection of image forgeries. *IEEE Transactions on Image Processing*, 2019. 1, 2, 5
- [2] Z. J. Barad and M. M. Goswami. Image forgery detection using deep learning: A survey. In 2020 6th International Conference on Advanced Computing and Communication Systems (ICACCS), pages 571–576, 2020. 1
- [3] B. Bayar and M. C. Stamm. Constrained convolutional neural networks: A new approach towards general purpose image manipulation detection. *IEEE Transactions on Information Forensics and Security*, 2018. 2
- [4] T. Aaron Gulliver & Saif alZahir Belal Ahmed. Image splicing detection using mask-rcnn. *Signal, Image and Video Processing*, 2020. 2
- [5] X. Bi, Y. Wei, B. Xiao, and W. Li. Rru-net: The ringed residual u-net for image splicing forgery detection. In 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pages 30–39, 2019.
- [6] Gajanan K. Birajdar and Vijay H. Mankar. Digital image forgery detection using passive techniques: A survey. *Digital Investigation*, 10(3):226 – 245, 2013. 1
- [7] L. Bondi, S. Lameri, D. Güera, P. Bestagini, E. J. Delp, and S. Tubaro. Tampering detection and localization through clustering of camera-based cnn features. In 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pages 1855–1864, 2017. 2
- [8] Yue Cao, Jiarui Xu, Stephen Lin, Fangyun Wei, and Han Hu. Gcnet: Non-local networks meet squeeze-excitation networks and beyond, 2019. 3, 4, 6
- [9] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs, 2017. 2
- [10] Xinru Chen, Chengbo Dong, Jiaqi Ji, Juan Cao, and Xirong Li. Image manipulation detection by multi-view multi-scale supervision, 2021. 3
- [11] Giovanni Chierchia, Giovanni Poggi, Carlo Sansone, and Luisa Verdoliva. Prnu-based forgery detection with regularity constraints and global optimization. In 2013 IEEE 15th International Workshop on Multimedia Signal Processing (MMSP), 2013. 2
- [12] Giovanni Chierchia, Giovanni Poggi, Carlo Sansone, and Luisa Verdoliva. A bayesian-mrf approach for prnu-based image forgery detection. *IEEE Transactions on Information Forensics and Security*, 9, 2014. 2
- [13] Davide Cozzolino and Luisa Verdoliva. Noiseprint: a cnnbased camera model fingerprint, 2018. 1, 2
- [14] Hao Dang, Feng Liu, Joel Stehouwer, Xiaoming Liu, and Anil K. Jain. On the detection of digital face manipulation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2020. 3

- [15] Jing Dong, Wei Wang, and Tieniu Tan. Casia image tampering detection evaluation database. pages 422–426, 07 2013.
 2, 5
- [16] P. Ferrara, T. Bianchi, A. De Rosa, and A. Piva. Image forgery localization via fine-grained analysis of cfa artifacts. *IEEE Transactions on Information Forensics and Security*, 7(5):1566–1577, 2012. 7
- [17] Aurobrata Ghosh, Zheng Zhong, Terrance E Boult, and Maneesh Singh. Spliceradar: A learned method for blind image forensics, 2019. 1
- [18] Thomas Gloe and Rainer Böhme. The 'dresden image database' for benchmarking digital image forensics. Association for Computing Machinery, 2010. 5
- [19] H. Guan, M. Kozak, E. Robertson, Y. Lee, A. N. Yates, A. Delgado, D. Zhou, T. Kheyrkhah, J. Smith, and J. Fiscus. Mfc datasets: Large-scale benchmark datasets for media forensic challenge evaluation. In 2019 IEEE Winter Applications of Computer Vision Workshops (WACVW), pages 63–72, 2019. 2, 5
- [20] Jie Hu, Li Shen, Samuel Albanie, Gang Sun, and Enhua Wu. Squeeze-and-excitation networks, 2019. 3, 4
- [21] Xuefeng Hu, Zhihan Zhang, Zhenye Jiang, Syomantak Chaudhuri, Zhenheng Yang, and Ram Nevatia. Span: Spatial pyramid attention network for image manipulation localization. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*. Springer International Publishing, 2020. 1, 2, 3, 5, 7
- [22] Iryna Korshunova, Wenzhe Shi, Joni Dambre, and Lucas Theis. Fast face-swap using convolutional neural networks, 2016. 1
- [23] N. Krawetz. A picture 's worth . . . digital image analysis and forensics version 2. 2007. 7
- [24] Jingyuan Li, Fengxiang He, Lefei Zhang, Bo Du, and Dacheng Tao. Progressive reconstruction of visual structure for image inpainting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 1
- [25] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection, 2018. 5
- [26] Huajian Liu, Martin Steinebach, and Kathrin Schölei. Improved manipulation detection with convolutional neural network for jpeg images. pages 1–6, 08 2019. 2
- [27] Xiaohong Liu, Yaojie Liu, Jun Chen, and Xiaoming Liu. Pscc-net: Progressive spatio-channel correlation network for image manipulation detection and localization, 2021. 2, 5, 6, 7
- [28] Babak Mahdian and Stanislav Saic. Using noise inconsistencies for blind image forensics. *Image Vision Comput.*, 27(10), 2009. 7
- [29] Gaël MAHFOUDI, Badr TAJINI, Florent RETRAINT, Frédéric MORAIN-NICOLIER, Jean Luc DUGELAY, and Marc PIC. Defacto: Image and face manipulation dataset. In 2019 27th European Signal Processing Conference (EU-SIPCO), 2019. 5
- [30] Ghazal Mazaheri, Niluthpol Chowdhury Mithun, Jawadul H. Bappy, and Amit K. Roy-Chowdhury. A skip connection architecture for localization of image manipulations. In Pro-

ceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, June 2019. 3

- [31] Aniruddha Mazumdar, Jaya Singh, Yosha Singh Tomar, and Prabin Kumar Bora. Universal image manipulation detection using deep siamese convolutional neural network, 2018. 2
- [32] G. Muzaffer and G. Ulutas. A new deep learning-based method to detection of copy-move forgery in digital images. In 2019 Scientific Meeting on Electrical-Electronics Biomedical Engineering and Computer Science (EBBT), 2019. 2
- [33] A. Novozámský, B. Mahdian, and S. Saic. Imd2020: A large-scale annotated dataset tailored for detecting manipulated images. In 2020 IEEE Winter Applications of Computer Vision Workshops (WACVW), pages 71–80, 2020. 2, 5
- [34] Ozan Oktay, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Mattias Heinrich, Kazunari Misawa, Kensaku Mori, Steven McDonagh, Nils Y Hammerla, Bernhard Kainz, Ben Glocker, and Daniel Rueckert. Attention u-net: Learning where to look for the pancreas, 2018. 3
- [35] Abdul Muntakim Rafi, Uday Kamal, Rakibul Hoque, Abid Abrar, Sowmitra Das, Robert Laganière, and Md. Kamrul Hasan. Application of densenet in camera model identification and post-processing detection, 2019. 2
- [36] Yuan Rao and Jiangqun Ni. A deep learning approach to detection of splicing and copy-move forgeries. 2016. 2
- [37] Abhijit Guha Roy, Nassir Navab, and Christian Wachinger. Recalibrating fully convolutional networks with spatial and channel 'squeeze & excitation' blocks, 2018. 5
- [38] Nikolay Sergievskiy and Alexander Ponamarev. Reduced focal loss: 1st place solution to xview object detection in satellite imagery, 2019. 6
- [39] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing, 2020. 1
- [40] Z. Shi, X. Shen, H. Kang, and Y. Lv. Image manipulation detection and localization based on the dual-domain convolutional neural networks. *IEEE Access*, 2018. 1
- [41] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition, 2015. 2
- [42] Xunyu Pan & Xing Zhang Siwei Lyu. Exposing region splicing forgeries with blind local noise estimation. *International Journal of Computer Vision*, 2014. 2
- [43] Carole H. Sudre, Wenqi Li, Tom Vercauteren, Sebastien Ourselin, and M. Jorge Cardoso. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. *Lecture Notes in Computer Science*, 2017. 5
- [44] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks, 2020. 2
- [45] Kunj Bihari Meena & Vipin Tyagi. A hybrid copy-move image forgery detection technique based on fourier-mellin and scale invariant feature transforms. *Multimedia Tools and Applications*, 2020. 2
- [46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. 2017. 2
- [47] Vinay Verma, Deepak Singh, and Nitin Khanna. Block-level double jpeg compression detection for image forgery localization, 2020. 2

- [48] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 3, 4
- [49] N. B. A. Warif, M. Y. I. Idris, A. W. A. Wahab, and R. Salleh. An evaluation of error level analysis in image forensics. In 2015 5th IEEE International Conference on System Engineering and Technology (ICSET), pages 23–28, 2015. 2
- [50] Liqing Wei, Wu, Dong, Zhang, and Sun. Developing an image manipulation detection algorithm based on edge detection and faster r-cnn. *Symmetry*, 11:1223, 10 2019. 2
- [51] B. Wen, Y. Zhu, R. Subramanian, T. Ng, X. Shen, and S. Winkler. Coverage — a novel database for copy-move forgery detection. In 2016 IEEE International Conference on Image Processing (ICIP), pages 161–165, 2016. 5
- [52] Yue Wu, Wael AbdAlmageed, and Prem Natarajan. Deep matching and validation network – an end-to-end solution to constrained image splicing localization and detection, 2017.
- [53] Yue Wu, Wael AbdAlmageed, and Prem Natarajan. Busternet: Detecting image copy-move forgery with source/target localization. In *European Conference on Computer Vision* (ECCV). Springer, 2018. 2
- [54] Y. Wu, W. AbdAlmageed, and P. Natarajan. 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), pages 9535–9544, 2019. 1, 2, 3, 5, 6, 7, 8
- [55] Lanqing Xue, Xiaopeng Li, and Nevin L. Zhang. Not all attention is needed: Gated attention network for sequence data, 2019. 3
- [56] Chao Yang, Huizhou Li, Fangting Lin, Bin Jiang, and Hao Zhao. Constrained r-cnn: A general image manipulation detection model, 2020. 2
- [57] Jiani Zhang, Xingjian Shi, Junyuan Xie, Hao Ma, Irwin King, and Dit-Yan Yeung. Gaan: Gated attention networks for learning on large and spatiotemporal graphs, 2018. 3
- [58] Peng Zhou, Bor-Chun Chen, Xintong Han, Mahyar Najibi, Abhinav Shrivastava, Ser Nam Lim, and Larry S. Davis. Generate, segment and refine: Towards generic manipulation segmentation, 2019. 2
- [59] Peng Zhou, Xintong Han, Vlad I. Morariu, and Larry S. Davis. Learning rich features for image manipulation detection, 2018. 2, 5, 6, 7
- [60] Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. Unet++: A nested u-net architecture for medical image segmentation, 2018. 2, 3