SeeTheSeams: Localized Detection of Seam Carving based Image Forgery in Satellite Imagery

Chandrakanth Gudavalli[†,1], Erik Rosten[†,1], Lakshmanan Nataraj [1], Shivkumar Chandrasekaran [1, 2], and B. S. Manjunath[1,2]

[†]Equal Contributors [1]Mayachitra, Inc.
[2]Electrical and Computer Engineering Department, UC Santa Barbara
Santa Barbara, California, USA

Abstract

Seam carving is a popular technique for content aware image retargeting. It can be used to deliberately manipulate images, for example, change the GPS locations of a building or displace/remove roads in a satellite image. This paper proposes a novel approach for detecting and localizing seams in such images. While there are methods to detect seam carving based manipulations, this is the first time that robust localization and detection of seam carving forgery is made possible. We also propose a seam localization score (SLS) metric to evaluate the effectiveness of localization. The proposed method is evaluated extensively on a large collection of images from different sources, demonstrating a high level of detection and localization performance across these datasets. The code and datasets curated during this work will be released to the public.

1. Introduction

Seam carving is a popular image manipulation technique [1, 3, 48] that is effective for content aware image retargeting. In seam carving, the input image is resized by removing or inserting "seams" which are defined as connected pixel paths from top-to-bottom or left-to-right. These seams are chosen by their optimality according to an energy function computed for each pixel. That is, the optimal seam is the seam with the lowest energy along its path. Commonly used energy functions in seam carving are computed by measuring the contrast of a pixel with its neighbors. Removing an optimal seam has fewer artifacts in resized images than a randomly chosen seam, and protects image content that is highly textured. Seam carving can be extended to remove entire objects from images by assigning the energy of object pixels to a low value such that the seams forcibly pass through the object marked for removal. Since seam carving leaves a large percentage of pixel values untampered, it poses a challenge to image forgery detection.

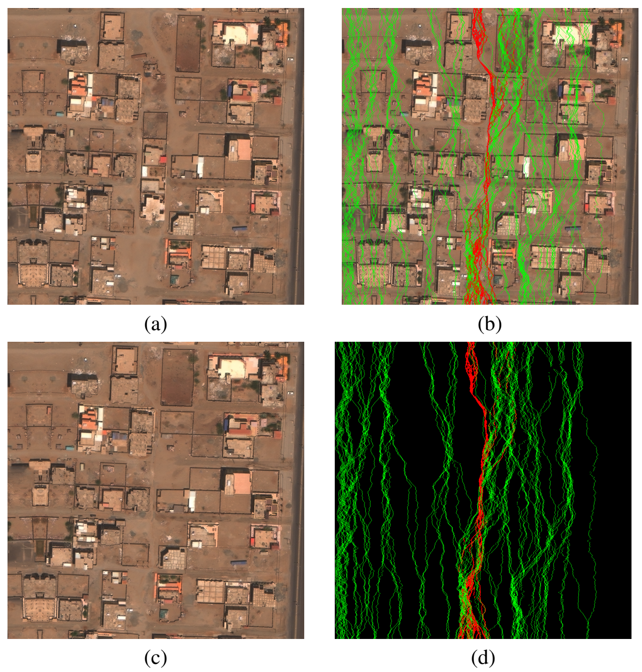

Figure 1. **(a)**. Pristine satellite image. **(b)**. Seam carved image with a few buildings removed overlaid with ground truth seam mask, where red seams indicate the pixels where seams are removed while green seams indicate the pixels where seams are inserted (red) seams. **(c)**. Seam carved image after removing a few objects (buildings). **(d)**. Seam prediction mask generated by our detector.

Seam carving is included as a feature in popular image editing software such as GIMP [33] and Photoshop [56]. The ease of access to these programs along with increasing availability of satellite image data from commercial satellites presents a growing problem for organizations that rely on accurate satellite data. Satellite images have been manipulated to influence public opinion such as in the Malaysia Airlines flight incident [29], nighttime flyovers of India during the festivals [12], and fake spliced images of bridges [15]. Furthermore, satellite imagery is particularly susceptible to seam carving based image manipulations. One reason is that in satellite imagery, objects of importance generally occupy fewer pixels than conventional images and consequently require fewer seams to remove, reducing the potential for visual artifacts in seam carved satellite images. In Figure 1 and Figure 2, we show an example of object removal where we remove a set of buildings from the original image and restore the original image dimensions. Moreover, objects can be "displaced" to change their geographical coordinates (latitude, longitude) using seam carving much more easily in satellite imagery than typical images. Since satellite images are captured from high altitudes, they often have large, smooth regions that are ideal for seams to pass through, making them good candidates for object displacement while preserving visual image quality. In Figure 4 , we show an example of object displacement accomplished by removing seams on one side of the building and inserting on the other. Displacing objects can be especially malicious in satellite imagery, where a single pixel can correspond with as much as $60m^2$ of ground area. In applications such as military target acquisition, this can be the difference between a successful and unsuccessful mission.

While there are several works proposed to detect seam carving based image forgery [14, 17, 24, 34, 41, 42, 44, 51, 59–62], to our knowledge, none of these methods have investigated localization of removed and inserted seams at the pixel level, and their design is restricted to classification at the image level, or classification and localization at a patch level. We develop and evaluate our method on satellite imagery as a case study due to the potential ramifications of seam carving based manipulations.

Towards addressing the above challenges, we propose a two stage, deep learning based seam carving detector with two key advantages over existing methods. The first advantage is the ability to *localize the seams* at pixel level resolution. This is invaluable in discovering the extent of potential manipulations in satellite imagery. By considering the location of seams, one may be able to discover not only that an object may have been removed in an image, but where that removed object used to be. The second advantage is *generalizability*. Since image level classifiers are trained on original satellite images, they are prone to become specialized to the distribution of data they are trained on. We demonstrate that our method is generalizable to not only training dataset distribution but different seam carving techniques. The main contributions of the paper are:

1. We propose a method for robust detection of seam carving manipulations and accurate localization of seams removed or inserted by seam carving in satellite imagery.

2. We develop the seam localization score (SLS), a specialized metric to better evaluate the localization performance of seam carving detectors on specific seams.

3. We have created unique, large seam carving datasets that we plan to release to the public [1].

The rest of this paper is organized as follows. In Section 2 we review existing forensic approaches related to our work. Section 3 details our proposed framework for localized detection of seam carving forgeries. Section 4 describes the datasets that are curated to carry out our seam carving forensic experiments. Metrics that are used to evaluate our models are described in Section 5. Experimental setups and results are detailed in Section 6. Finally, we conclude the paper in Section 7 by reviewing the pros and cons of our method and possible research areas that merit further exploration.

2. Related Work

Several works have been proposed to detect digital image manipulations (see [36, 54] for an overview). These works include the detection of specific image manipulations such as resampling [18, 46, 49], morphing [28, 45], copy-move [31], splicing [5, 50], seam carving [22, 32], and inpainting based object removal [57]. Several approaches also exploit JPEG blocking artifacts to detect tampered regions [16,35], while more recent efforts tend to exploit deep learning based approaches [8, 11, 43, 44].

Most image forensics techniques that have been developed so far target consumer images [6, 9, 10, 21, 39], which significantly differ in nature from satellite sensors (e.g. different compression schemes, color channels, orthorectification based post processing and so on). Furthermore, it has been observed that many of these techniques do not perform well when naively applied to overhead images [2, 7, 19, 40, 53, 58]. To address this issue, several forensic techniques that work for satellite/overhead imagery have recently been proposed [7, 26, 27, 58]. One method uses handcrafted watermarks to detect manipulations in satellite images [26]. Although this method is quite effective, it cannot be utilized if the watermark is not inserted at the time of image acquisition by a trustworthy source. Another technique leverages conditional GANs to detect and localize splicing forgeries in satellite images by estimating a forgery mask [58]. A second GAN based technique [27] encodes patches from an image into a low dimensional vector space that are used as input into a support vector machine (SVM) to detect the presence of forgeries at a patch level. Finally, Sat-SVDD [7] is a kernel-based one-class classification method that detects splicing forgery with the help of support vector data description (SVDD). In comparison to these methods, our paper explores the challenging case of detection and localization of seam carving based manipulations in satellite images.

[1] https://github.com/Mayachitra-Inc/SeeTheSeams

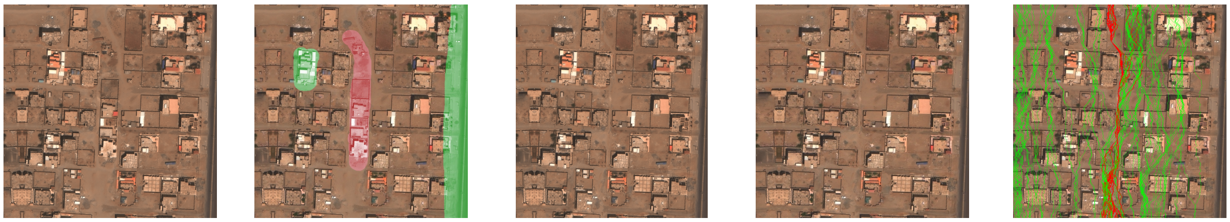

(a) (b) (c) (d) (e)

Figure 2. **(a)**. Original xView image **(b)**. Original image overlaid with masks on objects to be removed (red) and protected (green). **(c)**. Seam carved image with objects removed and the dimension of the image is reduced by 77 columns in the seam removal process. **(d)**. Final image after reinserting 77 seams into (c) to match dimensions of original image (a). **(e)**. Seam carved image overlaid with ground truth seam mask of inserted (green) seams and removal (red) seams.

There have been several works proposed over the past decade to reveal traces of seam carving in digital images [14, 17, 24, 34, 41, 42, 51, 59–62]. These include methods using steganalysis [51], hashing [17, 34], local binary patterns (LBP) [60–62], and deep learning based methods [14, 41, 42, 59]. In [24], the authors propose a method based on block artifacts to estimate the location of an object removed using seam carving. However, this method mainly focuses on finding a region in a seam carved image where an object has been removed and does not produce localization maps where seams have been removed/inserted. One deep learning method proposes a Convolutional Neural Network (CNN) based approach to perform image level classification using local binary patterns [14]. Another CNN based approach uses a customized network that learns and uses more effective features via joint optimization of feature extraction and pattern classification [59]. These two proposed methods perform image level classification to detect the presence of seam carving manipulation. One method that performs localization at a patch level employs a CNN called LFNet that is specifically designed to learn low-level features, capturing local artifacts from seam carving based image resizing [41]. The authors in [41] improved their network design with a new architecture, ILFNet [42]. The authors in [44] proposes a two stage model where stage-1 performs patch level localization and stage-2 performs image level classification. The method proposed in this paper localizes seams at a pixel level as opposed to the patch level strategies in prior work and focuses on satellite imagery, which is more resistant to seam carving artifacts than consumer images. In contrast to these proposed methods, our approach both detects images that are seam carved and also generates a seam localization mask which highlights pixels on the manipulated image where seams have been carved or inserted.

3. Proposed Approach

Here, we describe our approach for localized seam detection. At its core, the problem can be formulated as binary classification, where we would ultimately like to predict whether the given image is seam carved or not. In order to achieve this, we propose a two stage framework. In Section 3.1, we develop localization models that flag regions of the image that are seam carved. In Section 3.2, we perform a final classification whether the image has been manipulated by seam carving by fusing the outputs of stage 1.

3.1. Stage 1 - Localization of Seams

In order to obtain localized detection of seams in an image we implement a fully convolutional network to learn a mapping from satellite image to seam mask, where the seam mask contains the locations of removed or inserted seams as described in Section 4.1. U-Net has been extensively used in image segmentation [47], and its architecture is composed of a contracting path (encoder) followed by an expanding path (decoder). Encoder maps the input image to a feature vector, which will then be used by the decoder to flag seam carved regions of the input image. The basic encoder-decoder model's localization power is enhanced in U-Net by applying skip connections between encoder feature maps and decoder outputs. We chose EfficientNetB7 [52] as our baseline network since it performed the best when compared to other standard networks (see Section 6.1 for comparative experiments) EfficientNet is a baseline network composed of sequential mobile inverted bottleneck convolutions (MBConv) blocks that can be scaled up to improve accuracy at the cost of increased computation. This generates a family of models from EfficientNetB0 to EfficientNetB7 with the EfficientNetB7 model being the most accurate, but also computationally the most expensive. EfficientNetB7 also outperforms other encoder architectures when used for pixelwise classification, such as in [4]. For the decoder architecture, we use four sets of transposed convolutional, batch norm, dropout, and ReLU layers to upsample encoded features back to the size of input image. In Figure 3, we show an overview of our proposed method, where we see that the localization stage is comprised of two pixelwise classifiers - one model trained for seam removal detection and another model trained for seam insertion detection, both using U-Net and EfficientNetB7. On a test image, these two models output predicted seam localization masks flagging the regions where seams have been removed and inserted.

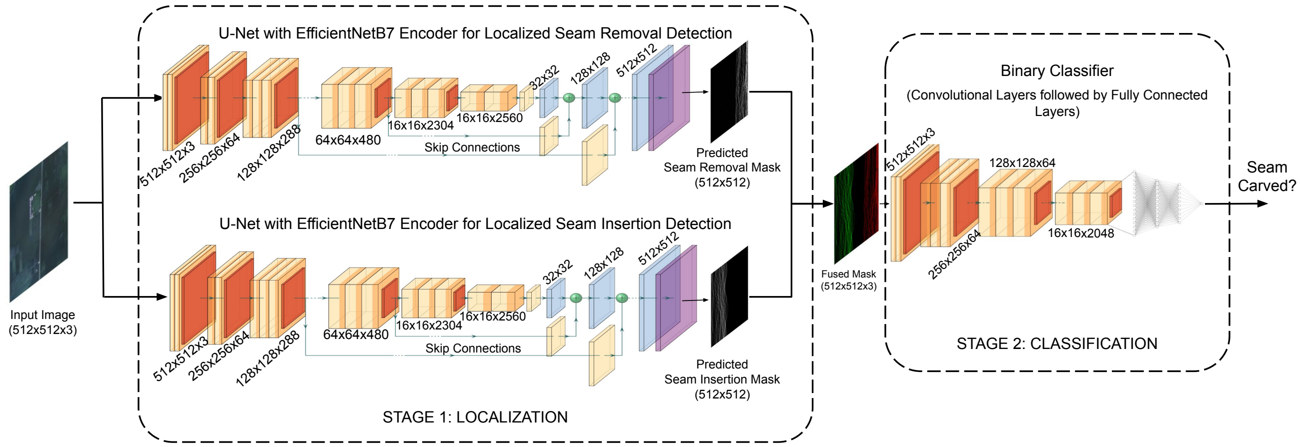

Figure 3. Overview of proposed two stage framework for localized seam carving detection.

3.2. Stage 2 - Classification

The pixelwise classifiers described above output two seam prediction masks with the same dimensions as the input image, one localizing removed seams and the other inserted seams. Both these masks are concatenated and fed as input to a standard CNN (ResNet50 [25]) to perform vanilla binary classification to identify if an image has been seam carved or not.

Standard CNN architectures are used as opposed to customized architectures in both the stages, to make the method simple and straightforward. Since the standard CNN architectures lead us to state of the art detection performance, having a specialized architecture for this problem seems unnecessary. However, one could think of coming up with an architecture that can minimize the computational complexity, instead of using beefy networks like EfficientNetB7, which is the future work for current methodology.

4. Datasets

In this section, we give a brief overview on the characteristics of source datasets used, and how we generated our forgery datasets. Three common satellite imagery datasets (xView [30], xBD [23], and Orbview-3 [20]) have been used to evaluate our method. The xView dataset contains $1,127$ images at varying high resolutions. Using these high resolution images, we generated a dataset of $53,943$ images by randomly cropping 512 x 512 regions from the original dataset. When dividing the images into a train-test-val split, all 512 x 512 images from a given xView sample are assigned to the same split. For xView, we allocated 70% of the images into training set, 15% into validation set, and remaining 15% into test set. The xBD dataset contains a total of $22,098$ pre-disaster and post-disaster RGB images of size 1024 x 1024, where we've preserved the original dataset's train-test-val split of 80:10:10 and randomly cropped 512 x 512 regions as with the xView dataset. The xBD dataset also contains ground truth building masks that can be used in seam carving object removal. These building masks are quite small and make up an average of 0.37% of all pixels in a particular training set image. On average, a building mask occupies 916 $pixels^2$. Since the xBD dataset contains satellite images at $0.3m$ per pixel, these building objects correspond with around $82m^2$ of ground area despite taking up a small number of pixels. In the most extreme cases, a building mask can take as little as $0.9m^2$ up to $15,943m^2$ of ground area. Finally, we curated a third dataset consisting of $48,000$ Orbview-3 images by randomly selecting non overlapping 512 x 512 regions. The train-test-val split is again maintained to be 80:10:10. While the xView and xBD datasets contain 8 bit RGB images (image intensities are in range 0 to 255), the Orbview-3 dataset has single channel 16 bit images. These three datasets provide variation in the geographical location of the images as well as different levels of color depth.

4.1. Ground Truth Seam Masks

Ground truth seam masks for training pixelwise classifiers are generated while seam carving original dataset images. As we remove seams from the original image, all the preserved pixel locations that were adjacent to removed seams are marked, indicating manipulation. As we insert seams into the image, all the pixel locations of the inserted seam are flagged. A visual example of a ground truth seam mask is shown overlaid on the resulting seam carved image in Figure 1b.

4.2. Pixelwise Classification Datasets

Datasets for pixelwise classification are generated by seam carving pristine samples to remove the top 10% of optimal seams and inserting seams to restore the original image dimension. All of our models are trained on 512 x 512 images, cropped from top left of the seam carved images and saved as PNG files. Ground truth seam masks are generated as described in Section 4.1 and similarly cropped. For the remainder of the paper if unspecified, forward energy is used to define the optimality of seams and examine the generalization capabilities of our models across several seam carving variants in Section 6.2.

4.3. Stage 2 Classification Datasets

Since we train models from each stage independently, stage 2 datasets are easily generated once pixelwise classifier training is complete. We obtain the prediction masks from trained model inference on cropped seam carved and original images to form a dataset of manipulated and pristine samples required for the image level classification.

5. Evaluation Metrics

In this section, we briefly describe the evaluation metrics used to select and assess our models. Since we're operating in a binary classification setting at a pixel level in the localization stage and at image level in the classification stage, we use metrics based on confusion matrices. The confusion matrix and the accuracy over the entire dataset is computed from cumulative confusion matrices for every sample. However, since seam carving based manipulations tend to remove less than 10 percent of the pixels from source datasets, pixelwise accuracy is an inadequate representation of the performance of our models as a naive method predicting all negatives will achieve above 90% accuracy. To address the inherent imbalance in our generated datasets, we prioritize several more relevant confusion matrix derived metrics like Precision, Recall, and F1-Score.

While useful, these three metrics are inherently biased towards the positive class, and independent of the number of true negatives. For seam carving localization, we would like to not only incorporate how close the predicted seams are to the true seams, but measure the efficacy of our model on correctly identifying untampered regions. One metric that satisfies this specification is the Matthews Correlation Coefficient (MCC), a balanced measure ranging between ± 1 that can be used regardless of the degree of class imbalance in a dataset due to its invariance to the choice of which class is positive or negative [13]. Detectors are trained on 512 x 512 patches cropped from the top left of seam carved images, and it should be noted that in cases where cropping resulted in all ground truth negatives for a particular sample, all metrics aside from pixelwise accuracy are set to 0 to avoid division by zero errors.

5.1. Customizing Confusion Matrix Metrics

While these confusion matrix metrics are widely used to evaluate binary classification performance, we make a slight adjustment to the way we calculate confusion matrices (and their derived metrics) for our specific use-case. Since seams are only one-pixel wide and our confusion matrices are calculated on a pixel-wise basis, they are extremely sensitive to the spatial distribution of the prediction mask. For example, take the seam insertion mask from Figure 4b. If we shift all the seams one pixel to the right and compare the shifted mask to the original, our confusion matrix metrics

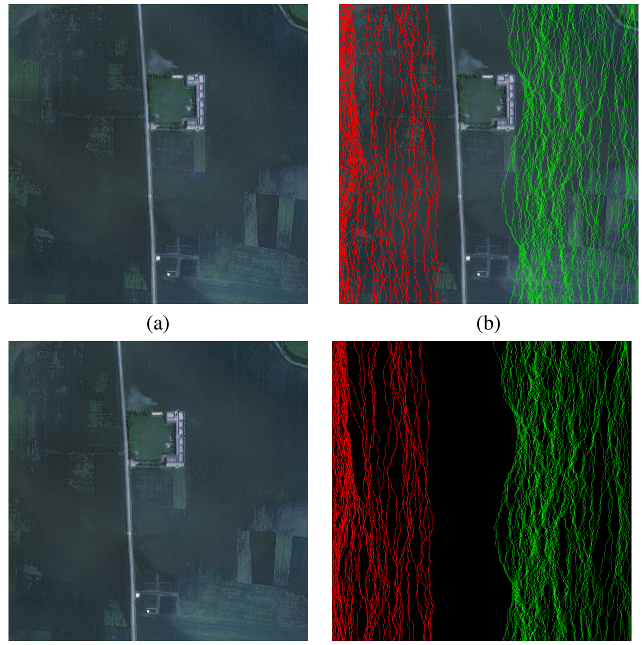

(a) (b)

(c) (d)

Figure 4. **(a)**. Pristine satellite image marked for "object displacement". **(b)**. Seam carved image overlaid with ground truth seam mask with removal seams marked in red and inserted seams marked in green. **(c)**. Seam carved image indicating "object displacement" where an entire strip in the center of the image has been displaced by a few pixels to the left. **(d)**. Prediction mask generated by our detector highlighting the areas where seams have been removed and inserted.

break down and report poor results. Specifically, the recall between the original and shifted masks becomes 0.194 when the shifted version is in fact localizing the seams quite well.

To properly represent the performance of our model's predictions, we modify the confusion matrix calculation such that we assign a true positive if our prediction is within a buffer of p pixels of a ground truth positive. This relaxation is similar to a metric used in evaluating road detectors, especially in the context of aerial imagery [38,55]. It's worth noting that if we assign a true positive in this way, we do not double count the ground truth positive used to evaluate predicted negative pixels. If we predict negatively and the ground truth at that location is positive but has been used, we assign it as a true negative. We employ the strategy described here with $p = 1$, and for the particular example in this section the recall between the original and shifted masks becomes 1.0, due the 1 pixel buffer described above. To make this clear, we refer to these customized metrics as MCC-1, F1 Score-1, Precision-1, and Recall-1 throughout the remainder of this paper.

5.2. Seam Localization Score (SLS)

To fully capture the localization performance of our models, we develop a metric based on the seams that are inserted or removed. By keeping track of the seams associated with manipulated pixels, we can evaluate how well our model localizes each specific seam. For a particular vertical seam s taken from an image of height h and width w, the corresponding ground truth seam mask will associate h pixels with seam s. Then, we sum over each seam pixel from $1 : h$, the absolute distance to the nearest predicted positive in each row. If we do not predict a positive pixel in a particular row, the absolute distance is set to w. Finally, we normalize by the number of rows h. We refer to this metric the seam localization score (SLS). For an image with N seams we compute an image level SLS by summing the score for each seam, and dividing by N. The SLS for any particular seam ranges from 0, perfect overlap, to w. When the SLS for a particular seam is less than one, we can interpret that on average, the seam was less than one pixel away from its ground truth location.

6. Experimental Results

Here, we describe the experiments that are carried out towards localized seam carving detection. First, we cover our multistage approach towards localized detection of seam carving, using pixelwise classifiers, and then present the performance of localization models using the metrics detailed in Section 5. Then, we demonstrate the generalizability of our model across different training distributions and seam carving algorithms.

Encoder Architecture	F1Score-1	MCC-1	SLS
MobileNetV2	0.779	0.776	5.95
ResNet50	0.631	0.651	17.93
ResNet101	0.603	0.628	23.73
EfficientNetB7	**0.911**	**0.903**	**1.56**

Table 1. Performance of seam removal detectors with various encoder architectures, trained and tested on xView.

Dataset	MCC-1	F1Score-1
xView	0.956	0.959
xBD	0.938	0.935
OrbView-3	0.942	0.943

Table 2. Generalizability of a seam insertion pixelwise classifier trained on xView, and tested on all three datasets.

	Tested on xView	Tested on xBD	Tested on Orbview-3
Trained on xView	0.903	0.821	0.711
Trained on xBD	0.722	0.894	0.708
Trained on Orbview-3	0.721	0.726	0.889

Table 3. MCC-1 scores of EfficientNetB7 seam removal pixelwise classifiers, trained and tested on different datasets.

6.1. Pixelwise Classification

To achieve more fine-grained localization of seams and and incorporate the spatial relationships of input images, we train a step-down, step-up fully convolutional neural network based on U-Net as described in Section 3.1 to detect removed and inserted seams at pixel-level resolution (Figure 1). We report stage 1 evaluation metrics from Section 5 using a variety of encoder architectures in Table 1 on our xView test dataset, as generated in Section 4.2. The best performing model used an EfficientNetB7 encoder architecture, achieving a MCC-1 score of 0.903. We note that the performance increase gained by using EfficientNet cannot be solely due to the increase in model capacity. The second best performing network was MobileNetV2, which has the smallest number of trainable parameters among all networks tested. Our seam insertion detector of the same EfficientNetB7 architecture, achieves a MCC-1 score of 0.956 on the xView test set. We also observe that although trained on xView images, the seam insertion detector generalizes well to xBD and Orbview-3 datasets with an MCC-1 of 0.938 and 0.942 respectively.

6.2. Generalizability of Pixelwise Classifiers

Here, we summarize several experiments demonstrating the generalizability of pixelwise classifiers across datasets and various seam carving techniques. Pixelwise classifiers trained on our xView dataset work well on xBD and Orbview-3 datasets and vice versa. Moreover, even though pixelwise classifiers are trained with datasets forged using forward energy, they are generalizable to different seam carving techniques.

6.2.1 Generalizability Across Datasets

Table 2 shows that pixelwise classifiers for seam insertion detection are generalizable across different datasets. A seam insertion detector trained on xView has only minor drops in performance on xBD and Orbview-3 test sets.

We also observe that while seam removal detectors are not as generalizable as seam insertion detectors, they still perform well on different datasets as shown in Table 3. Although the decrease in MCC-1 scores for seam removal detectors are larger when stress testing across datasets, we note that the scores themselves are still quite good, and are indicative of adequate performance for stage 2 classification.

6.2.2 Generalizability Across Seam Carving Methods

We also tested the generalizability of our models across various seam carving methods. So far, all of our results have been reported on datasets that have been generated using forward energy seam carving. We report test evaluation metrics on our xView dataset in Table 4, where we have generated test sets using backward energy, frequency tuned

Seam carving techniques	Evaluation Metrics			
	Precision-1	Recall-1	F1Score-1	MCC-1
Backward Energy [3]	0.936	0.834	0.882	0.874
Forward Energy [48]	0.956	0.869	0.911	0.903
Frequency Tuned Saliency Map Method [1]	0.817	0.425	0.559	0.578
Seam Merging Method [37]	0.862	0.521	0.649	0.658

Table 4. Seam carving technique generalizability of the seam removal pixelwise classifier, trained on forward energy seam carved xView images and tested on xView images seam carved using other techniques.

Dataset	xView	xBD	xBD_OR	Orbview-3
Accuracy	99.29	99.08	99.73	98.46

Table 5. Test set accuracy(%) of stage 2 binary classifiers.

saliency maps, and seam merging variations of seam carving.In this table, the model has been trained on only forward energy seam carving data, but we see good performance across different seam carving techniques, with the lowest performance on the dataset generated using saliency map based seam carving. This is most likely due to how different the seams from saliency map seam carving look compared to the other techniques.

6.3. Stage 2 Classification

In stage 2, we generate a dataset of stage 1 predictions on manipulated and pristine images to train a final binary classifier to check if the input image has been seam carved. We use a stage 1 model trained on the xBD dataset as described in 4.2. This stage 1 model is used to obtain predictions on a combination of xBD datasets. The first xBD dataset that we predict on is described in 4.2. We remove the 10% most optimal seams, which are often distributed throughout the image, and reinsert seams to restore the original image dimensions. The second xBD dataset we predict on incorporates the original dataset's ground truth building masks and uses seam carving to remove a building and reinsert seams to restore the original image dimensions. We call this dataset "xBD_OR". Thus, our stage 2 model is trained on a combination of stage 1 predictions on best seam removed and object removed images. Finally, we generate best seam removed test datasets of xView, xBD and Orbview-3 using individually trained pixelwise classifiers. For example, the Orbview-3 test dataset for stage 2 is generated using an Orbview-3 trained pixelwise classifier. Although the xView and Orbview-3 test datasets are generated using different pixelwise classifiers than the training dataset, our stage 2 model performs very well on both, achieving almost 99% accuracy. This shows that the stage 2 final classification model is robust to the pixelwise classification model used to generate the input prediction masks.

6.4. Seam Carving Retargetting Ratios

In this section, we present results of two experiments where we vary the number of seams removed and inserted from the original xView dataset.

In Table 6, we provide MCC-1 and SLS scores for both seam removal and insertion detectors as well as stage 2 binary classification accuracies on test datasets of varying seam carving retargetting ratios using a model that is trained on 10% seam carved data. Test datasets are generated by seam carving 512x512 patches from xview test split by different percentages. We can see from this table that our seam carving detector is remarkably generalizable to other seam carving retargetting ratios despite being trained on only 10% seam carved data. In terms of an overall image classification, our framework achieves over 99% accuracy at detecting seam carved images on all retargetting ratios except 2%, where it achieves the lowest accuracy at 98.56%. In general, the SLS score of the seam removal detector decreases the further away we move from 10% seam carving, while the MCC-1 only decreases as we increase the retargetting ratio. This showcases the usefulness of the SLS metric as a seam carving localization metric. In the case of 2% seam carved data our seam removal detector achieves its best MCC-1 score of 0.918 due to the large amount of non manipulated pixels in the ground truth and predicted masks. However, the SLS score of 2.05 shows that our model is not as good at predicting seam locations as the MCC-1 score might lead us to expect. In the seam insertion case, we see good performance across the range of retargetting ratios tested. The lowest MCC-1 score for seam insertion is 0.920 on 50% seam carved data, which is higher than the best MCC-1 score for seam removal at 0.918 on 2% seam carved data. The seam insertion SLS scores show that our model is able to localize inserted seams within 1 pixel precision across all retargetting ratios. Although we show results on a model that was trained using a 10% seam carving ratio in this paper, the results were similar for models trained on other percentages of seam carving too.

In Table 7, we show results of training on datasets of different seam carving retargetting ratios. We include the results on the 10% seam carved test set, and provide evaluation metrics on 20%, 30%, 40%, and 50% seam carved test datasets. This table shows that our method is applicable to other seam carving retargetting ratios and achieves similar results. The SLS score is around 1 for all values tested, and the F1 Score-1 and MCC-1 scores are high. Notably, the 20% seam carving dataset has the highest performance while the 50% seam carved dataset has the worst.

SCR (%)	SR Detector		SI Detector		Stage2
	MCC-1	SLS	MCC-1	SLS	Acc
2	0.918	2.05	0.943	0.26	98.56
4	0.913	1.48	0.951	0.14	99.02
6	0.910	1.30	0.953	0.12	99.06
8	0.908	1.18	0.954	0.12	99.14
10	0.903	1.56	0.959	0.11	99.29
20	0.886	1.05	0.953	0.14	99.23
30	0.831	1.51	0.945	0.25	99.34
40	0.697	4.48	0.934	0.42	99.51
50	0.471	21.79	0.920	0.61	99.51

Table 6. Performance of Stage 1 and Stage 2 models, trained on a 10% seam carved images from the xView dataset and evaluated on xView test sets with varying numbers of removed and inserted seams(also referred as Seam Carving Ratio (SCR)).

Seam Carving Ratio	F1Score-1	MCC-1	SLS
10%	0.911	0.903	1.56
20%	0.986	0.893	0.838
30%	0.907	0.891	0.776
40%	0.901	0.883	0.892
50%	0.865	0.847	1.23

Table 7. Performance of Stage 1 models trained on xview datasets generated by inserting and removing different percentage of seams.

6.5. Performance on post processed images

We explored the performance of our models on the dataset of images that are JPEG compressed after seam carving. Retraining the models on post processed datasets resulted in the similar image level detection, while the models that are not trained on post processed images has shown drop in performance, as shown in Table 8. Same trend has been observed when we replaced JPEG compression with rotation as post processing step, as shown in Table 9.

JPEG Compression Quality Factor	Model trained without JPEG	Model trained with JPEG
60	51.96	84.86
70	53.88	90.11
80	59.06	95.06
90	74.19	97.81
No Comp	99.26	99.18

Table 8. Stage-2 Test accuracy(%): Improvement in detection accuracy when the model is trained on post processed images.

6.6. Comparison against Existing Approaches

We compare our method with two seam carving detection methods, that uses LBP [60] and ILFNet [42]. Although these methods are not specifically designed for localized detection in satellite imagery, we compare our image level stage 2 results with them as they are performing image level classification to detect seam carved images. Comparative results are shown in Table 10, where all of the forensic techniques have been trained on the xView training set generated as in Section 4.2. We observe that although ILFNet achieves a comparable accuracy with the proposed method on the xView test set, it's generalizability performance drops around 5% and 2% when testing on xBD and Orbview-3 while our method drops less than 1%. Similarly, a method that utilizes local binary pattern based feature extraction combined with an SVM classifier performs reasonably well on the xView test set, but fails to generalize to other datasets.

Rotation (Degrees)	Model trained without rotated images	Model trained with rotated images
45	52.11	91.26
60	56.63	96.78
75	63.14	97.66
90	99.26	99.11
0	99.26	99.11

Table 9. Stage-2 Test accuracy(%): Improvement in detection accuracy when the model is trained on post processed images.

Forensic Technique	xView	xBD	Orbview-3
LBP based detection [60]	91.72	82.71	78.41
ILFNet [42]	99.03	94.96	96.19
Proposed method	**99.29**	**98.86**	**98.43**

Table 10. Test accuracy(%): Comparison with other image level classifiers, trained on xView, and tested on all datasets.

7. Conclusion

In this paper, we proposed a method to detect and localize seam carving based manipulations in satellite images. We use a two stage approach that first localizes removed/inserted seams via pixelwise classification and then performs a final classification if an image has been seam carved. We enable localization of seams as well as a generalizable framework across different datasets and seam carving techniques. Finally, we detailed the curation of three unique, large seam carving satellite image datasets that will be released to the public. Although the proposed method is not specifically restricted to satellite imagery, we present our findings on satellite images as a case study and leave further evaluation on more conventional images with varying compression schemes and preprocessing to be explored in future work.

Acknowledgment

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA), the National Geospatial-Intelligence Agency (NGA) and the Air Force Research Laboratory (AFRL) under the contract number FA8750-16-C-0078. The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

References

[1] R. Achanta and S. Süsstrunk. Saliency detection for content-aware image resizing. In *2009 16th IEEE International Conference on Image Processing (ICIP)*, pages 1005–1008, 2009. 1, 7

[2] Luqman Ali, Teerasit Kasetkasem, Faisal Ghaffar Khan, Thitiporn Chanwimaluang, and Hiroki Nakahara. Identification of inpainted satellite images using evalutionary artificial neural network (eann) and k-nearest neighbor (knn) algorithm. In *2017 8th International Conference of Information and Communication Technology for Embedded Systems (IC-ICTES)*, pages 1–6. IEEE, 2017. 2

[3] Shai Avidan and Ariel Shamir. Seam carving for content-aware image resizing. *ACM Trans. Graph.*, 26(3):10, 2007. 1, 7

[4] Bhakti Baheti, Shubham Innani, Suhas Gajre, and Sanjay Talbar. Eff-unet: A novel architecture for semantic segmentation in unstructured environment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 358–359, 2020. 3

[5] Jawadul H Bappy, Amit K Roy-Chowdhury, Jason Bunk, Lakshmanan Nataraj, and BS Manjunath. Exploiting spatial structure for localizing manipulated image regions. In *ICCV*, 2017. 2

[6] 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*, 28(7):3286–3300, 2019. 2

[7] Emily R Bartusiak, Sri Kalyan Yarlagadda, David Güera, Paolo Bestagini, Stefano Tubaro, Fengqing M Zhu, and Edward J Delp. Splicing detection and localization in satellite imagery using conditional gans. In *2019 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*, pages 91–96. IEEE, 2019. 2

[8] Belhassen Bayar and Matthew C Stamm. A deep learning approach to universal image manipulation detection using a new convolutional layer. In *Proceedings of the 4th ACM Workshop on Information Hiding and Multimedia Security*, pages 5–10, 2016. 2

[9] Luca Bondi, Luca Baroffio, David Güera, Paolo Bestagini, Edward J Delp, and Stefano Tubaro. First steps toward camera model identification with convolutional neural networks. *IEEE Signal Processing Letters*, 24(3):259–263, 2016. 2

[10] Luca Bondi, Silvia Lameri, David Guera, Paolo Bestagini, Edward J Delp, Stefano Tubaro, et al. Tampering detection and localization through clustering of camera-based cnn features. In *CVPR Workshops*, pages 1855–1864, 2017. 2

[11] Jason Bunk, Jawadul H Bappy, Tajuddin Manhar Mohammed, Lakshmanan Nataraj, Arjuna Flenner, BS Manjunath, Shivkumar Chandrasekaran, Amit K Roy-Chowdhury, and Lawrence Peterson. Detection and localization of image forgeries using resampling features and deep learning. *arXiv preprint arXiv:1707.00433*, 2017. 2

[12] Deborah Byrd. Fake image of diwali still circulating. https://earthsky.org/earth/fake-image-of-india-during-diwali-versus-the-real-thing, 2018. Accessed: 2021-03-16. 1

[13] Davide Chicco and Giuseppe Jurman. The advantages of the matthews correlation coefficient (mcc) over f1 score and accuracy in binary classification evaluation. *BMC Genomics*, 21, 2020. 5

[14] Luiz Fernandoda Silva Cieslak, Kelton Augustopontara Da Costa, and Joao PauloPapa. Seam carving detection using convolutional neural networks. In *2018 IEEE 12th International Symposium on Applied Computational Intelligence and Informatics (SACI)*, pages 000195–000200. Ieee, 2018. 2, 3

[15] Jane Edwards. Nga's todd myers: China uses gan technique to tamper with earth images. https://www.executivegov.com/2019/04/ngas-todd-myers-china-uses-gan-technique-to-tamper-with-earth-images/, 2019. Accessed: 2021-03-16. 2

[16] Hany Farid. Exposing digital forgeries from jpeg ghosts. *IEEE transactions on information forensics and security*, 4(1):154–160, 2009. 2

[17] Wei Fei, Yang Gaobo, Li Leida, Xia Ming, and Zhang Dengyong. Detection of seam carving-based video retargeting using forensics hash. *Security and Communication Networks*, 8(12):2102–2113, 2015. 2, 3

[18] Arjuna Flenner, Lawrence Peterson, Jason Bunk, Tajuddin Manhar Mohammed, Lakshmanan Nataraj, and Bangalore S Manjunath. Resampling forgery detection using deep learning and a-contrario analysis. *Electronic Imaging*, 2018(7):212–1, 2018. 2

[19] Mohamed Mahmoud Fouad, Eslam Magdy Mostafa, and Mohamed Abdelmoneim Elshafey. Detection and localization enhancement for satellite images with small forgeries using modified gan-based cnn structure. *International Journal of Advances in Intelligent Informatics*, 6(3):278–289, 2020. 2

[20] GeoEye. Usgs eros archive - commercial satellites - orbview 3. `https://doi.org/10.5066/F7J38R0R`, (accessed Oct 24, 2019). 4

[21] Michael Goebel, Lakshmanan Nataraj, Tejaswi Nanjundaswamy, Tajuddin Manhar Mohammed, Shivkumar Chandrasekaran, and BS Manjunath. Detection, attribution and localization of gan generated images. *Electronic Imaging*, 2021(4):276–1, 2021. 2

[22] Qingmei Gong, Qingqing Shan, Yongzhen Ke, and Jing Guo. Detecting the location of seam and recovering image for seam inserted image. *Journal of Computational Methods in Sciences and Engineering*, 18(2):499–509, 2018. 2

[23] Ritwik Gupta, Richard Hosfelt, Sandra Sajeev, Nirav Patel, Bryce Goodman, Jigar Doshi, Eric T. Heim, Howie Choset, and Matthew E. Gaston. xbd: A dataset for assessing building damage from satellite imagery. *CoRR*, abs/1911.09296, 2019. 4

[24] Rong Han, Yongzhen Ke, Ling Du, Fan Qin, and Jing Guo. Exploring the location of object deleted by seam-carving. *Expert Systems with Applications*, 95:162–171, 2018. 2, 3

[25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. 4

[26] A. T. S. Ho, Xunzhan Zhu, and W. M. Woon. A semi-fragile pinned sine transform watermarking system for content authentication of satellite images. In *Proceedings. 2005 IEEE International Geoscience and Remote Sensing Symposium, 2005. IGARSS '05.*, volume 2, pages 4 pp.–, 2005. 2

[27] János Horváth, David Güera, Sri Kalyan Yarlagadda, Paolo Bestagini, Fengqing Maggie Zhu, Stefano Tubaro, and Edward J Delp. Anomaly-based manipulation detection in satellite images. *networks*, 29:21, 2019. 2

[28] Sabah Jassim and Aras Asaad. Automatic detection of image morphing by topology-based analysis. In *2018 26th European Signal Processing Conference (EUSIPCO)*, pages 1007–1011. IEEE, 2018. 2

[29] Andrew Kramer. Russian images of malaysia airlines flight 17 were altered, report finds. https://www.nytimes.com/2016/07/16/world/europe/malaysia-airlines-flight-17-russia.html, 2016. Accessed: 2021-03-16. 1

[30] Darius Lam, Richard Kuzma, Kevin McGee, Samuel Dooley, Michael Laielli, Matthew Klaric, Yaroslav Bulatov, and Brendan McCord. xview: Objects in context in overhead imagery. *CoRR*, abs/1802.07856, 2018. 4

[31] Jian Li, Xiaolong Li, Bin Yang, and Xingming Sun. Segmentation-based image copy-move forgery detection scheme. *IEEE Transactions on Information Forensics and Security*, 10(3):507–518, 2015. 2

[32] Yanan Li, Ming Xia, Xin Liu, and Gaobo Yang. Identification of various image retargeting techniques using hybrid features. *Journal of Information Security and Applications*, 51:102459, 2020. 2

[33] Liquid Rescale GIMP plugin. http://liquidrescale.wikidot.com/en:examples. Accessed: 2021-03-16. 1

[34] Wenjun Lu and Min Wu. Seam carving estimation using forensic hash. In *13th ACM multimedia workshop on Multimedia and security*, pages 9–14, 2011. 2, 3

[35] Weiqi Luo, Jiwu Huang, and Guoping Qiu. Jpeg error analysis and its applications to digital image forensics. *IEEE Transactions on Information Forensics and Security*, 5(3):480–491, 2010. 2

[36] Babak Mahdian and Stanislav Saic. A bibliography on blind methods for identifying image forgery. *Signal Processing: Image Communication*, 25(6):389–399, 2010. 2

[37] Kazu Mishiba and Masaaki Ikehara. Image resizing using improved seam merging. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1261–1264. IEEE, 2012. 7

[38] Volodymyr Mnih and Geoffrey E Hinton. Learning to detect roads in high-resolution aerial images. In *European Conference on Computer Vision*, pages 210–223. Springer, 2010. 5

[39] Tajuddin Manhar Mohammed, Jason Bunk, Lakshmanan Nataraj, Jawadul H Bappy, Arjuna Flenner, BS Manjunath, Shivkumar Chandrasekaran, Amit K Roy-Chowdhury, and Lawrence A Peterson. Boosting image forgery detection using resampling features and copy-move analysis. *Electronic Imaging*, 2018(7):118–1, 2018. 2

[40] Daniel Mas Montserrat, János Horváth, SK Yarlagadda, Fengqing Zhu, and Edward J Delp. Generative autoregressive ensembles for satellite imagery manipulation detection. *arXiv preprint arXiv:2010.03758*, 2020. 2

[41] Seung-Hun Nam, Wonhyuk Ahn, Seung-Min Mun, Jinseok Park, Dongkyu Kim, In-Jae Yu, and Heung-Kyu Lee. Content-aware image resizing detection using deep neural network. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 106–110. IEEE, 2019. 2, 3

[42] Seung-Hun Nam, Wonhyuk Ahn, In-Jae Yu, Myung-Joon Kwon, M. Son, and H. Lee. Deep convolutional neural network for identifying seam-carving forgery. *ArXiv*, abs/2007.02393, 2020. 2, 3, 8

[43] Lakshmanan Nataraj, Michael Goebel, Tajuddin Manhar Mohammed, Shivkumar Chandrasekaran, and BS Manjunath. Holistic image manipulation detection using pixel cooccurrence matrices. *Electronic Imaging*, 2021(4):277–1, 2021. 2

[44] Lakshmanan Nataraj, Chandrakanth Gudavalli, Tajuddin Manhar Mohammed, Shivkumar Chandrasekaran, and BS Manjunath. Seam carving detection and localization using two-stage deep neural networks. In *Machine Learning, Deep Learning and Computational Intelligence for Wireless Communication*, pages 381–394. Springer, 2021. 2, 3

[45] Tom Neubert. Face morphing detection: An approach based on image degradation analysis. In *International Workshop on Digital Watermarking*, pages 93–106. Springer, 2017. 2

[46] Alin C Popescu and Hany Farid. Exposing digital forgeries by detecting traces of resampling. *IEEE Transactions on signal processing*, 53(2):758–767, 2005. 2

[47] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *CoRR*, abs/1505.04597, 2015. 3

[48] Michael Rubinstein, Ariel Shamir, and Shai Avidan. Improved seam carving for video retargeting. *ACM Transactions on Graphics (SIGGRAPH)*, 27(3):1–9, 2008. 1, 7

[49] Seung-Jin Ryu and Heung-Kyu Lee. Estimation of linear transformation by analyzing the periodicity of interpolation. *Pattern Recognition Letters*, 36:89–99, 2014. 2

[50] Ronald Salloum, Yuzhuo Ren, and C-C Jay Kuo. Image splicing localization using a multi-task fully convolutional network (mfcn). *Journal of Visual Communication and Image Representation*, 51:201–209, 2018. 2

[51] Anindya Sarkar, Lakshmanan Nataraj, and Bangalore S Manjunath. Detection of seam carving and localization of seam insertions in digital images. In *11th ACM workshop on Multimedia and security*, pages 107–116. ACM, 2009. 2, 3

[52] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *CoRR*, abs/1905.11946, 2019. 3

[53] Adam Van Etten. You only look twice: Rapid multi-scale object detection in satellite imagery. *arXiv preprint arXiv:1805.09512*, 2018. 2

[54] Savita Walia and Krishan Kumar. Digital image forgery detection: a systematic scrutiny. *Australian Journal of Forensic Sciences*, pages 1–39, 2018. 2

[55] Jan D Wegner, Javier A Montoya-Zegarra, and Konrad Schindler. A higher-order crf model for road network extraction. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1698–1705, 2013. 5

[56] What's new in adobe photoshop cs4 - photoshop 11. http://www.photoshopsupport.com/photoshop-cs4/what-is-new-in-photoshop-cs4.html. Accessed: 2021-03-16. 1

[57] Qiong Wu, Shao-Jie Sun, Wei Zhu, Guo-Hui Li, and Dan Tu. Detection of digital doctoring in exemplar-based inpainted images. In *Machine Learning and Cybernetics, 2008 International Conference on*, volume 3, pages 1222–1226. IEEE, 2008. 2

[58] Sri Kalyan Yarlagadda, David Güera, Paolo Bestagini, Fengqing Maggie Zhu, Stefano Tubaro, and Edward J Delp. Satellite image forgery detection and localization using gan and one-class classifier. *Electronic Imaging*, 2018(7):214–1, 2018. 2

[59] Jingyu Ye, Yuxi Shi, Guanshuo Xu, and Yun-Qing Shi. A convolutional neural network based seam carving detection scheme for uncompressed digital images. In *International Workshop on Digital Watermarking*, pages 3–13. Springer, 2018. 2, 3

[60] Ting Yin, Gaobo Yang, Leida Li, Dengyong Zhang, and Xingming Sun. Detecting seam carving based image resizing using local binary patterns. *computers & security*, 55:130–141, 2015. 2, 3, 8

[61] Dengyong Zhang, Qingguo Li, Gaobo Yang, Leida Li, and Xingming Sun. Detection of image seam carving by using weber local descriptor and local binary patterns. *Journal of information security and applications*, 36:135–144, 2017. 2, 3

[62] Dengyong Zhang, Gaobo Yang, Feng Li, Jin Wang, and Arun Kumar Sangaiah. Detecting seam carved images using uniform local binary patterns. *Multimedia Tools and Applications*, 79(13):8415–8430, 2020. 2, 3