

# Manipulation Detection in Satellite Images Using Deep Belief Networks

János Horváth, Daniel Mas Montserrat, Hanxiang Hao, Edward J. Delp

Video and Image Processing Laboratory (VIPER)  
School of Electrical Engineering  
Purdue University  
West Lafayette, Indiana, USA

## Abstract

*Satellite images are more accessible with the increase of commercial satellites being orbited. These images are used in a wide range of applications including agricultural management, meteorological prediction, damage assessment from natural disasters and cartography. Image manipulation tools including both manual editing tools and automated techniques can be easily used to tamper and modify satellite imagery. One type of manipulation that we examine in this paper is the splice attack where a region from one image (or the same image) is inserted (“spliced”) into an image. In this paper, we present a one-class detection method based on deep belief networks (DBN) for splicing detection and localization without using any prior knowledge of the manipulations. We evaluate the performance of our approach and show that it provides good detection and localization accuracies in small forgeries compared to other approaches.*

## 1. Introduction

Satellite images can be used in many applications including meteorological measurements, such as precipitation prediction [1], thunderstorm detection [2] and wind speed and direction estimation [3]. The analysis of satellite images can also be an efficient way to assess regional infrastructure levels [4, 5], classify crops in agricultural applications [6, 7], forest characterization [8], scene classification [9, 10] or to estimate soil moisture [11, 12].

The number of commercial satellites is increasing exponentially [13] with many of these platforms having advanced imaging sensors. These satellites have provided a large number of image datasets available to the public [14, 15, 16], such as Planet Labs or the European Space Agency image datasets [17, 18].

Satellite images can be easily forged and manipulated.




Figure 1. Examples of manipulated images (left), manipulation masks (center), and manipulation heatmap (right).

Image and video editing tools such as GIMP [19] or Photoshop [20] can be used to create realistic forgeries. Machine learning techniques [21] can quickly manipulate images in an automatic way, without manual edition. The easy access to image manipulation methods, coupled with an increasing amount of digital content, poses a problem for institutions that rely on satellite imagery. Some examples where satellite images were manipulated in order to bias public opinion include the Malaysia Airlines Flight incident [22], the nighttime flyovers of India during the Diwali festivals [23] and the fake Chinese bridge spliced images [24].

Common manipulations include splicing (cropping and pasting regions from the same or other image sources) and machine learning-based forgeries, typically generated with Generative Adversarial Networks (GANs) [21]. While several methods that validate the authenticity and integrity of images and videos have been proposed [25, 26, 27], verification of a satellite image authenticity still remains an unsolved problem. The wide range of forgery objects




Figure 2. Method overview for generating the heatmap and the detection score

and techniques makes their detection challenging. Common image forensic methods are often developed for images captured by consumer cameras and fail on satellite imagery as they differ greatly in compression schemes, post-processing, sensors, and color channels. There is a need for manipulation detection methods that are accurate regardless of the tampering objects and technologies used for image capturing.

In this paper, we present a method based on Deep Belief Networks (DBN) [28] to detect and localize splicing forgeries in satellite images. Figure 1 shows the results of the proposed method. The left column shows the manipulated images, the middle column contains the ground truth images of the manipulation and the right column shows the detection results with the accurate locations and shapes of the spliced objects. A deep belief network is a probabilistic generative model that is composed of multiple layers of Restricted Boltzmann Machines (RBM) [29]. The network learns the distribution of the training data in an unsupervised manner and can be used to detect out-of-distribution data. The overview of the proposed method is shown in Figure 2. We first generate a set of image patches cropped from the original satellite image. The deep belief network takes the patches as input and outputs the reconstructed patches. In this paper, we use two stacked Restricted Boltzmann Machines (RBMs) for the deep belief network. At the next step, an error map is computed by combining the difference of the original and reconstructed patches. The final reconstructed heatmap is obtained by normalizing the error map. We evaluate our method using a dataset containing splicing forgeries, described in Section 3. Additionally, we show that our method can be used as a One-Class classifier and we evaluate it with the MNIST [30] dataset.

The main contributions of the paper are that we present a new technique to generate datasets containing satellite images with splicing manipulations. We also introduce a method for splicing manipulation detection and localization using DBNs that does not require any manipulated data during training. We then evaluate our method with multiple

configurations of RBMs.

## 2. Related Work

The forensics community has described many techniques to detect various types of forgeries and manipulations of images. Some of these techniques include detecting tampering by finding double-JPEG compression artifacts [31], using neural networks with domain adaptation [32] or using saturation cues [33]. Other techniques focus on detecting splicing in images [34, 35]. The work presented in [34], proposes a new feature-based algorithm to detect splicing in images without any prior information. They extract expressive features that capture traces left locally by in-camera processing in order to detect manipulations. The work in [35] proposes a method to extract the camera model fingerprint in which the scene content is largely suppressed and model-related artifacts are enhanced. We compare our method performance with both of the previously described splicing detection methods. Most of these methods fail when used with satellite imagery. Because the image acquisition process differs between “cameras” and satellites, methods designed for the former do not transfer properly to the latter. These differences include satellite sensor technologies and post-processing steps such as orthorectification, radiometric corrections, and compression.

Several techniques to analyze the integrity of satellite images have been presented using hand-crafted features [36] and data-driven approaches including both supervised [37] and unsupervised [38, 39] methods. Manipulated images are required to use for the supervised techniques during training, but not for unsupervised methods. Watermarking techniques have been used for tampering detection [36]. The authors in [37] introduce a supervised approach to detect and localize manipulation in satellite images based on conditional GAN [40]. It is trained on original and manipulated images in order to map an input image to a forgery mask. The work in [38] presents a GAN-based method to encode patches extracted from the input image into a low

dimensional feature vector. Then a one-class support vector machine (SVM) is used to detect if a patch contains forgeries by comparing the distribution of the feature learned from the original image patches. The Sat-SVDD method [39] follows a similar approach by using a modified Support Vector Data Description (SVDD) [41] to detect splicing forgeries. Sat-SVDD is a kernel-based one-class classification method that performs minimum volume estimation with a neural network to extract the statistical regularities of the dataset. First, the SVDD encodes each image patch to a latent space. Then, it aims to minimize the distance between each patch and a predefined center point in the latent space. By doing so, the model forces the latent vector of original images inside a hypersphere. At testing, the latent vectors outside the hypersphere are considered as forged data.

In this paper, we use deep belief networks (DBN) [42]. Deep belief networks are composed of stacked layers of restricted Boltzmann machines. [43]. Boltzmann machines are based on Hopfield networks [44, 45] which are recurrent neural networks where each input node (or “visible unit”) is symmetrically interconnected. These networks can learn the data distribution from a limited set of training examples and reconstruct noisy or incomplete samples through an optimization process of “energy” minimization. Boltzmann machines follow the same structure as Hopfield networks but add a hidden layer connected to the visible layer. The size and computation of these networks increase exponentially concerning the data dimensionality because all nodes are interconnected. In order to reduce such complexity, [46] introduces a training process based on contrastive divergence minimization. Furthermore, [28] presents the restricted Boltzmann machine (RBM) that reduces the network complexity and training time by removing the connections between nodes (or units) of the same layer. Many variations of the restricted Boltzmann machine have been presented. Some examples include the higher-order Boltzmann machine [47], the conditional Boltzmann machine [45] and the mean field Boltzmann [48] machine. While regular Boltzmann machines are parametrized with a Bernoulli distribution, many different distributions, including Gaussian have been explored [29, 49, 50]. Furthermore, several training methodologies [42, 51, 52] has been presented to reduce the time complexity of the training process. More recent work combine DBNs with modern networks such as GANs [53].

### 3. Dataset

In this paper, we use satellite imagery from the region of Slovenia taken from the Sentinel program [54]. A total of 293 orthorectified images with an image resolution of  $1000 \times 1000$  pixels are collected. We use 100 of the 293 orthorectified images to create manipulated images. 19 different objects are spliced into the 100 images generating a total




Figure 3. Manipulated dataset examples: two manipulated images with spliced objects (left) and their respective ground truth masks (right).

of 500 manipulated images with their corresponding manipulation ground truth masks. The 19 objects include clouds, planes, smoke and drones images. Figure 3 shows some examples from the manipulated dataset. As shown in Figure 4, we splice the objects in different locations, rotation angles and sizes including  $16 \times 16$ ,  $32 \times 32$ ,  $64 \times 64$ ,  $128 \times 128$ , and  $256 \times 256$  pixels. Therefore, the forgery dataset contains a total of 693 images: 193 original images and 500 manipulated images (generated from 100 original images). We then split the dataset into training and testing sets. The training dataset consist of 98 original images from the original orthorectified images and the testing dataset consists of 595 images: 95 original images and 500 manipulated images with their corresponding ground truth masks.

## 4. Method

### 4.1. Splicing Detection

Our approach has three steps: (1) patch extraction and normalization, (2) patch reconstruction with DBN, and (3) manipulation heatmap and manipulation score estimation as shown in Figure 2.

Our method starts by cropping overlapping patches of  $64 \times 64$  pixels from the full-resolution input images. The image patches are obtained by splitting the full-resolution image into a set of overlapping patches with a stride of 32 and 8 pixels during training and testing respectively. We use a larger stride for training to reduce the computational time and a smaller stride during testing to have a more accurate estimate. Finally, each patch is normalized to be in the range of 0 in 1.

Then, the extracted patches are used as input to a DBN.

The DBN learns to reconstruct the input patch. We assume that patches from original (non-manipulated) images have different statistical properties than patches containing splicing manipulations. Therefore, we train the DBN only with original images. By doing so, the model will properly encode and reconstruct patches from original images but will fail with forged patches.

Finally, we compute the mean square error (MSE) between the input patch and the reconstructed patch. Patches from original (non-manipulated) images will be properly reconstructed, and therefore the reconstruction error will be small. However, patches with splicing manipulations will be reconstructed poorly leading to a higher MSE. We create a manipulation heatmap representing the probability of containing a splice manipulation by combining the MSE at each patch. We average the MSE properly to take into account the overlap within patches. Once the manipulation heatmap has been estimated, we follow the same approach as in [39] to produce a localization mask of an anomaly by thresholding.

For the anomaly detection task, we compute a detection score to determine the existence of anomaly given a manipulation heatmap. Following the work presented in [39], we compute a detection score as:

$$d(M) = \frac{\max(M) - \mu_M}{\sqrt{\frac{\sum_{x \in M} (x - \mu_M)^2}{\max(|M|)}}}, \quad (1)$$

Where  $M$  is the estimated manipulation heatmap composed by  $N$  pixels,  $\mu_M = \sum_{x \in M} \frac{x}{N}$  is the average mask value, and  $\max(M)$  is the maximum value of  $M$ . For a selected threshold  $T$ , our method considers an input image as manipulated if  $d(M) > T$  and original (non-manipulated) otherwise.

## 4.2. Network Architecture

We use a DBN composed by two stacked Restricted Boltzmann Machines (RBM). The first RBM takes the image patches as input and outputs a feature vector also, referred as hidden representation, of size 2916. The second RBM takes as input the hidden representation of the first layer and outputs a final hidden representation of the same size as input. By enforcing a hidden representation with lower dimensionality than the input patch, the network is forced to learn compact features that contain the statistical information of the original image which is used for the image reconstruction task. Each RBM is specified by the distributions used in their visible (input) and hidden layers. We examined three different settings for the RBMs: Gaussian-Bernoulli RBM [49], Gaussian-Gaussian RBM [55] and Uniform-Uniform RBM.

The Gaussian-Bernoulli RBM [49] has been used in previous work [55] and has been shown that it is able to learn

features that resemble those learned by the Independent Component Analysis (ICA) algorithm [56]. In this RBM the input data is modeled with a Gaussian distribution and the hidden layer is modeled by a Bernoulli distribution. The conditional probability of the visible layer, given the hidden layer, is defined by the following Normal distribution:

$$p(\mathbf{v}|\mathbf{h}) = \mathcal{N}(\mathbf{h}\mathbf{W}^T + \mathbf{a}, \mathbf{I}) \quad (2)$$

Where  $\mathbf{h}$  are the values of the hidden layer,  $\mathbf{W}$  the linear parameters,  $\mathbf{a}$  the bias term, and  $\mathbf{I}$  the identity matrix. The conditional probability of the hidden layer given the visible layer is defined by the following Bernoulli distribution:

$$p(\mathbf{h}|\mathbf{v}) = \prod_i \sigma(\mathbf{w}_i \mathbf{v} + b_i)^{h_i} (1 - \sigma(\mathbf{w}_i \mathbf{v} + b_i))^{1-h_i} \quad (3)$$

Where  $w_i$  and  $b_i$  are the  $i$ th row vector of the linear and the bias terms, and  $\sigma(\cdot)$  is the Sigmoid function. By modeling the input with a Normal distribution, it can represent data with a wide range of values (such as images). The Bernoulli distribution enforces a binary (and therefore sparse) representation of the hidden layer [57].

The Gaussian-Gaussian RBM [55] uses Normal distributions to model both visible and hidden layers. Such method has been shown to be able to learn features that resemble Principal Component Analysis (PCA) features [58]. The conditional probability of the visible layer, given the hidden layer, is represented in the same way as in Equation 2. The conditional probability of the hidden layer given the visible layer is:

$$p(\mathbf{h}|\mathbf{v}) = \mathcal{N}(\mathbf{W}\mathbf{v} + \mathbf{b}, \mathbf{I}) \quad (4)$$

Where  $\mathbf{v}$  are the values of the visible layer,  $\mathbf{W}$  the linear parameters,  $\mathbf{b}$  the bias term, and  $\mathbf{I}$  the identity matrix. By modelling both visible and hidden layer with a Normal distribution, the network is able to learn a dense representation of the images.

Additionally, we explore using another setting, the Uniform-Uniform RBM. In this network, both visible and hidden layer are modeled with uniform distributions. The conditional probability of the visible layer, given the hidden layer is:

$$p(\mathbf{v}|\mathbf{h}) = U(\lambda_1(\mathbf{h}\mathbf{W}^T + \mathbf{a}), \lambda_2(\mathbf{h}\mathbf{W}^T + \mathbf{a})) \quad (5)$$

The conditional probability of the hidden layer given the visible layer is:

$$p(\mathbf{h}|\mathbf{v}) = U(\lambda_1(\mathbf{W}\mathbf{v} + \mathbf{b}), \lambda_2(\mathbf{W}\mathbf{v} + \mathbf{b})) \quad (6)$$

Where  $\mathbf{v}$  are the values of the visible layer,  $\mathbf{h}$  are the values of the hidden layer,  $\mathbf{W}$  the linear parameters, and  $\mathbf{a}$  and  $\mathbf{b}$  the bias terms. We empirically select the values  $\lambda_1$  and  $\lambda_2$  that provide a low reconstruction error on the training set. In this work we use  $\lambda_1 = \frac{3}{4}$  and  $\lambda_2 = \frac{5}{4}$ .




Figure 4. An overview of how the manipulated image dataset is created

### 4.3. One-Class Classifier

The proposed method can be used for the one-class classification task. In a one-class classification problem, the classifier learns important features to identify a known target class based on a training set containing only images of such class. During testing, the classifier distinguishes between images of the target class and images from other unknown previously unseen classes, based on the features learned during training. Such an approach is especially useful for outlier and anomaly detection.

In order to evaluate our method as a one-class classifier, we train and test it to perform as handwritten digit recognition with the MNIST dataset [30]. In this scenario, the network only has access to images of one target digit during training. Then, the network has to distinguish between the target digit and other digits never seen during training. When performing a one-class classification, we do not divide the image within patches but use the DBN to reconstruct the complete input image. In this setting, the manipulation heatmap  $M$  and the manipulation score  $d(M)$  are not estimated. Instead, the reconstruction error is directly used to classify an image as the known class or not. If the reconstruction error is below a chosen threshold  $T$ , the image is classified as the known class. Section 5 presents some experiments of one-class classification with MNIST dataset [30].

## 5. Experimental Results

We evaluate our method as a splicing detection and localization method with the dataset presented in Section 3. We also examine the proposed method of the one-class classification task with the MNIST [30] dataset that contains images of handwritten digits from 0-9.

We train and test our DBN based on Uniform-Uniform RBMs (UU-DBN) with the complete training splicing manipulation datasets and compare it with previous work. Note that the UU-DBN is trained only with original images and no manipulated images were used during training. We compute the ROC curves for detection and Precision/Recall curves for localization tasks by changing the

threshold used to the estimated manipulation mask and manipulation score. Note that when evaluating the localization performance, there is a larger number of original than manipulated pixels. Therefore we use the AUC of the Precision/Recall curve, which is a balanced metric, to evaluate the localization performance. Table 1 presents our results compared with previous methods. Different ROC and Precision/Recall are shown for each of the different sizes of the splice objects used. In other words,  $ROC_{16}$  is the ROC for manipulated images with spliced objects of size  $16 \times 16$ . We can observe that UU-DBN provides similar or better results compared to previous methods. It is able to correctly detect and localize small splicing forgeries. The method provides a better P/R localization scores, specifically for forgeries with sizes from  $16 \times 16$  pixels to  $128 \times 128$  pixels. Figure 5 provides some visual examples of the estimated manipulation heatmaps for each of the methods. The estimated manipulation heatmaps show that the method is able to properly distinguish between splicing manipulations and the background image.

Additionally, we train and test four different DBNs as one-class classifiers with MNIST images. We use DBNs based on RBMs with Bernoulli-Bernoulli (BB-DBN), Gaussian-Bernoulli (GB-DBN), Gaussian-Gaussian (GG-DBN) and Uniform-Uniform (UU-DBN) distributions. We compare our results with previous one-class classification methods. Table 2 shows the AUC scores for each method. We can observe that the presented method provides competitive results with previous methods and that a DBN based on Uniform-Uniform distribution provides higher accuracy. The Gaussian-Bernoulli, Gaussian-Gaussian and Uniform-Uniform DBNs provide a similar AUC score while the Bernoulli-Bernoulli DBN provides a lower score. The binary representation used in the input and hidden layers of the Bernoulli-Bernoulli DBN provides features less flexible that in turn provide worse classification performance.

## 6. Conclusions

Satellite image manipulation detection is a challenging task due to the wide range of manipulations that can be present and the large variety of imaging technology used in






Figure 5. From rows left to right: input images, ground truth masks, Cozzolino *et al* [34], Yarlagadda *et al* [38], Horváth *et al* [39], Cozzolino *et al* [35] and UU-DBN.

Table 1. AUC scores (%) for the detection and localization task (ROC and P/R metrics). The subscript denotes the manipulation size. Best performing methods are **bold**.

Detection					
	Cozzolino <i>et al</i> [34]	Yarlagadda <i>et al</i> [38]	Horváth <i>et al</i> [39]	Cozzolino <i>et al</i> [35]	UU-DBN
ROC <sub>16</sub>	49.7	50.7	45.3	47.7	<b>58.2</b>
ROC <sub>32</sub>	50.4	59.6	57.7	47.2	<b>68.5</b>
ROC <sub>64</sub>	68.6	75.9	80.0	50.8	<b>82.6</b>
ROC <sub>128</sub>	84.8	81.5	87.4	56.7	<b>88.3</b>
ROC <sub>256</sub>	86.2	83.8	<b>89.9</b>	55.4	89.6

Localization					
	Cozzolino <i>et al</i> [34]	Yarlagadda <i>et al</i> [38]	Horváth <i>et al</i> [39]	Cozzolino <i>et al</i> [35]	UU-DBN
P/R <sub>16</sub>	0.0	0.0	0.1	0.0	<b>7.5</b>
P/R <sub>32</sub>	0.5	0.3	1.4	0.1	<b>13.3</b>
P/R <sub>64</sub>	7.8	2.5	18.1	2.5	<b>31.7</b>
P/R <sub>128</sub>	31.2	18.3	34.4	4.6	<b>40.5</b>
P/R <sub>256</sub>	48.5	37.8	<b>55.7</b>	7.8	48.8

orbiting satellites. In this paper we present an unsupervised splicing manipulation and one-class classification method based on deep belief networks. Without any prior knowledge from the manipulation information during the training process, the method provides competitive results compared to the previous work, especially with small splicing manipulations. We also evaluate multiple configurations of our network in a one-class classification framework providing competitive results compared to the common one-class classification methods. To improve the performance of our proposed method we plan to examine more different structures of Deep Belief Networks and Restricted Boltzmann Machines. Furthermore, we plan to examine other type of generative methods such as variational autoencoders.

## 7. Acknowledgment

This material is based on research sponsored by DARPA and Air Force Research Laboratory (AFRL) under agreement number FA8750-16-2-0173. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA and Air Force Research Laboratory (AFRL) or the U.S. Government.

Address all correspondence to Edward J. Delp, ace@ecn.purdue.edu .

Table 2. AUC scores (%) for the detection task Receiver operating characteristic (ROC), The best performing DBN configuration is marked in blue. The methods with best performance are bold.

Class	BB-DBN	GB-DBN	GG-DBN	UU-DBN	OCSVM /SVDD [59]	DCAE [60]	ONE-CLASS DEEP SVDD [61]	RCAE [62]
0	97.0 ± 0.1	99.1 ± 0.0	99.1 ± 0.0	<b>99.3 ± 0.1</b>	97.0 ± 0.7	97.6 ± 0.7	97.2 ± 1.2	<b>99.8±0.0</b>
1	99.7 ± 0.0	<b>99.8±0.0</b>	<b>99.8±0.0</b>	<b>99.8±0.1</b>	98.6 ± 0.8	98.3 ± 0.6	97.1 ± 1.5	<b>99.8±0.0</b>
2	79.9 ± 0.3	86.4 ± 0.2	86.5 ± 0.1	<b>89.5 ± 0.1</b>	78.9 ± 2.1	85.4 ± 2.4	97.3 ± 0.8	<b>98.8±0.6</b>
3	85.2 ± 0.2	91.3 ± 0.1	91.4 ± 0.1	<b>93.5 ± 0.1</b>	84.3 ± 2.6	86.7 ± 0.9	97.3 ± 1.0	<b>99.3±0.0</b>
4	90.0 ± 0.2	95.1 ± 0.1	95.2 ± 0.1	<b>96.1 ± 0.0</b>	93.6 ± 1.5	86.5 ± 2.0	97.3 ± 1.1	<b>99.2±0.0</b>
5	86.2 ± 0.2	92.9 ± 0.1	93.0 ± 0.1	<b>94.7 ± 0.1</b>	74.6 ± 4.5	78.2 ± 2.7	97.1 ± 1.2	<b>99.2±0.0</b>
6	94.9 ± 0.1	98.3 ± 0.0	98.3 ± 0.0	<b>98.7 ± 0.1</b>	95.4 ± 1.2	94.6 ± 0.5	96.7 ± 1.3	<b>99.8±1.0</b>
7	93.8 ± 0.1	95.6 ± 0.1	95.8 ± 0.1	<b>96.4 ± 0.0</b>	91.9 ± 1.5	92.3 ± 1.0	97.4 ± 1.0	<b>99.2±0.1</b>
8	78.4 ± 0.1	84.4 ± 0.1	85.1 ± 0.1	<b>87.4 ± 0.1</b>	87.9 ± 1.5	86.5 ± 1.6	97.2 ± 0.6	<b>98.5±2.0</b>
9	90.0 ± 0.1	94.5 ± 0.1	94.5 ± 0.1	<b>95.3 ± 0.1</b>	93.3 ± 1.2	90.4 ± 1.8	96.6 ± 1.5	<b>99.0±1.3</b>

## References

- [1] V. Lebedev, V. Ivashkin, I. Rudenko, A. Ganshin, A. Molchanov, S. Ovcharenko, R. Grokhovetskiy, I. Bushmarinov, and D. Solomentsev, "Precipitation nowcasting with satellite imagery," *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, August 2019, Anchorage, AK.
- [2] Y. Zhang, S. Wistar, J. Li, M. Steinberg, and J. Wang, "Storm detection by visual learning using satellite images," *Proceedings of the IEEE Transactions on Geoscience and Remote Sensing*, March 2016, Yokohama, Japan.
- [3] I. Sahoo, J. Guinness, and B. Reich, "Estimating atmospheric motion winds from satellite image data using space-time drift models," *arXiv:1902.09653*, February 2019.
- [4] P. K. Suraj, A. Gupta, M. Sharma, S. B. Paul, and S. Banerjee, "On monitoring development using high resolution satellite images," *arXiv:1712.02282*, December 2017.
- [5] B. Oshri, A. Hu, P. Adelson, X. Chen, P. Dupas, J. Weinstein, M. Burke, D. Lobell, and S. Ermon, "Infrastructure quality assessment in africa using satellite imagery and deep learning," *Proceedings of the International Conference on Knowledge Discovery & Data Mining*, August 2018, London, United Kingdom.
- [6] M. Rußwurm, S. Lefèvre, and M. Körner, "Breizhcrops: A satellite time series dataset for crop type identification," *Proceedings of the International Conference on Machine Learning Time Series Workshop*, June 2019, Long Beach, CA.
- [7] J. Brandt, "Spatio-temporal crop classification of low-resolution satellite imagery with capsule layers and distributed attention," *arXiv:1904.10130*, April 2019.
- [8] A. Chauve, C. Vega, S. Durrieu, F. Bretar, T. Allouis, M. P. Deseilligny, and W. Puech, "Advanced full-waveform lidar data echo detection: Assessing quality of derived terrain and tree height models in an alpine coniferous forest," *International Journal of Remote Sensing*, vol. 30, no. 19, pp. 5211–5228, September 2009.
- [9] A. Davari, V. Christlein, S. Vesal, A. Maier, and C. Riess, "GMM Supervectors for Limited Training Data in Hyperspectral Remote Sensing Image Classification," *Proceedings of the International Conference on Computer Analysis of Images and Patterns*, pp. 296–306, July 2017, Ystad, Sweden.
- [10] M. Shimoni, D. Borghys, R. Heremans, C. Perneel, and M. Acheroy, "Land-cover classification using fused polar and polinsar features," *Proceedings of the European Conference on Synthetic Aperture Radar*, pp. 1–4, June 2008, Friedrichshafen, Germany.
- [11] N. Efremova, D. Zausaev, and G. Antipov, "Prediction of soil moisture content based on satellite data and sequence-to-sequence networks," *Proceedings of the Conference on Neural Information Processing Systems Women in Machine Learning Workshop*, December 2018, Montreal, Canada.
- [12] D. D. Alexakis, F. K. Mexis, A. K. Vozinaki, I. N. Daliakopoulos, and I. K. Tsanis, "Soil moisture content estimation based on sentinel-1 and auxiliary earth observation products. A hydrological approach," *Sensors*, June 2017.
- [13] "UCS satellite database," <https://www.ucsusa.org/resources/satellite-database/>.
- [14] G. Xia, X. Bai, J. Ding, Z. Zhu, S. Belongie, J. Luo, M. Datcu, M. Pelillo, and L. Zhang, "Dota: A large-scale dataset for object detection in aerial images," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3974–3983, June 2018, Salt Lake City, UT.
- [15] S. M. Azimi, C. Henry, L. Sommer, A. Schumann, and E. Vig, "Skyscapes fine-grained semantic understanding of aerial scenes," *Proceedings of the IEEE International Conference on Computer Vision*, October 2019, Seoul, Korea.
- [16] Y. Yang and S. Newsam, "Bag-of-visual-words and spatial extensions for land-use classification," *Proceedings of the ACM SIGSPATIAL International Symposium on Advances in Geographic Information Systems*, pp. 270–279, November 2009, San Jose, CA.
- [17] R. Gupta, R. Hosfelt, S. Sajeev, N. Patel, B. Goodman, J. Doshi, E. Heim, H. Choset, and M. Gaston, "xBD:

- A dataset for assessing building damage from satellite imagery,” *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 10–17, June 2019.
- [18] M. Schmitt, L. H. Hughes, C. Qiu, and X. X. Zhu, “SEN12MS – a curated dataset of georeferenced multi-spectral sentinel-1/2 imagery for deep learning and data fusion,” *arXiv:1906.07789*, June 2019.
- [19] The GIMP Development Team, “GIMP,” <https://www.gimp.org>.
- [20] E. Bailey, *Adobe Photoshop: A Beginners Guide to Photoshop Lightroom - The 52 Photoshop Lightroom Tricks You Didn't Know Existed!*, vol. 1, CreateSpace Independent Publishing Platform, 2016, North Charleston, SC.
- [21] S. Nam, Y. Kim, and S. J. Kim, “Text-adaptive generative adversarial networks: Manipulating images with natural language,” *Proceedings of the Advances in Neural Information Processing Systems*, December 2018, Montreal, Canada.
- [22] A. E. 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>.
- [23] D. Byrd, “Fake image of diwali still circulating,” <https://earthsky.org/earth/fake-image-of-india-during-diwali-versus-the-real-thing>.
- [24] J. Edwards, “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/>.
- [25] A. Rocha, W. Scheirer, T. Boult, and S. Goldenstein, “Vision of the unseen: Current trends and challenges in digital image and video forensics,” *ACM Computing Surveys*, vol. 43, no. 4, October 2011.
- [26] V. Schetinger, M. Iuliani, A. Piva, and M. Oliveira, “Image forgery detection confronts image composition,” *Computers & Graphics*, vol. 68, no. C, pp. 152–163, November 2017.
- [27] F. Bartolini, A. Tefas, M. Barni, and I. Pitas, “Image authentication techniques for surveillance applications,” *Proceedings of the IEEE*, vol. 89, no. 10, pp. 1403–1418, October 2001.
- [28] G. E. Hinton and R. R. Salakhutdinov, “Reducing the dimensionality of data with neural networks,” *Science*, vol. 313, no. 5786, pp. 504–507, July 2006.
- [29] Y. Freund and D. Haussler, “Unsupervised learning of distributions on binary vectors using two layer networks,” *Proceedings of the Advances in Neural Information Processing Systems*, December 1994, Denver, CO.
- [30] Y. LeCun and C. Cortes, “MNIST handwritten digit database,” <http://yann.lecun.com/exdb/mnist/>, January 2010.
- [31] M. Barni, A. Costanzo, and L. Sabatini, “Identification of cut&paste tampering by means of double-JPEG detection and image segmentation,” *Proceedings of the IEEE International Symposium on Circuits and Systems*, pp. 1687–1690, May 2010, Paris, France.
- [32] D. Cozzolino, J. Thies, A. Rössler, C. Riess, M. Nießner, and L. Verdoliva, “Forensictransfer: Weakly-supervised domain adaptation for forgery detection,” *arXiv:1812.02510*, December 2018.
- [33] S. McCloskey and M. Albright, “Detecting GAN-Generated Imagery Using Saturation Cues,” *Proceedings of the IEEE International Conference on Image Processing*, pp. 4584–4588, September 2019, Taipei, Taiwan.
- [34] D. Cozzolino, G. Poggi, and L. Verdoliva, “Splicebuster: A new blind image splicing detector,” *Proceedings of the IEEE International Workshop on Information Forensics and Security*, pp. 1–6, November 2015, Rome, Italy.
- [35] D. Cozzolino and L. Verdoliva, “Noiseprint: A cnn-based camera model fingerprint,” *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 144–159, 2020.
- [36] A. T. S. Ho and W. M. Woon, “A semi-fragile pinned sine transform watermarking system for content authentication of satellite images,” *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium*, vol. 2, pp. 1–4, July 2005, Seoul, South Korea.
- [37] E. R. Bartusiak, S. Kalyan Yarlagadda, D. Güera, F. M. Zhu, P. Bestagini, S. Tubaro, and E. J. Delp, “Splicing detection and localization in satellite imagery using conditional gans,” *Proceedings of the IEEE International Conference on Multimedia Information Processing and Retrieval*, pp. 91–96, March 2019, San Jose, CA.
- [38] S. Kalyan Yarlagadda, D. Güera, P. Bestagini, F. Zhu, S. Tubaro, and E. Delp, “Satellite image forgery detection and localization using gan and one-class classifier,” *Proceedings of the IS&T International Symposium on Electronic Imaging*, vol. 2018, no. 7, pp. 214–214–9, February 2018, Burlingame, CA.
- [39] J. Horvath, D. Guera, S. Kalyan Yarlagadda, P. Bestagini, F. Maggie Zhu, S. Tubaro, and E. J. Delp, “Anomaly-based manipulation detection in satellite images,” *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 62–71, June 2019, Long Beach, CA.
- [40] P. Isola, J. Zhu, T. Zhou, and A. A. Efros, “Image-to-image translation with conditional adversarial networks,” *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5967–5976, July 2017, Honolulu, HI.
- [41] D. M. J. Tax and R. P. W. Duin, “Support vector data description,” *Machine Learning*, vol. 54, no. 1, pp. 45–66, January 2004.
- [42] G. E. Hinton, S. Osindero, and Y.-W. Teh, “A fast learning algorithm for deep belief nets,” *Neural Computer*, vol. 18, no. 7, July 2006.
- [43] P. Smolensky, *Information Processing in Dynamical Systems: Foundations of Harmony Theory*, MIT Press, Cambridge, MA, 1986.



- [44] J. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proceedings of the National Academy of Sciences*, vol. 79, no. 8, pp. 2554–2558, April 1982.
- [45] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, "A learning algorithm for boltzmann machines," *Cognitive Science*, vol. 9, pp. 147–169, January 1985.
- [46] G. E. Hinton, "Training products of experts by minimizing contrastive divergence," *Neural Computation*, vol. 14, no. 8, pp. 1771–1800, August 2002.
- [47] T. J. Sejnowski, "Higher-order boltzmann machines," *Proceedings of the Applied Inverse Problems Conference*, vol. 151, no. 1, pp. 398–403, 1986.
- [48] C. Peterson and J. R. Anderson, "A mean field theory learning algorithm for neural networks," *Complex Systems*, vol. 1, pp. 995–1019, 1987.
- [49] J. Chu, H. Wang, H. Meng, P. Jin, and T. Li, "Restricted boltzmann machines with gaussian visible units guided by pairwise constraints," *IEEE Transactions on Cybernetics*, vol. 49, no. 12, August 2019.
- [50] M. Welling, M. Rosen-Zvi, and G. Hinton, "Exponential family harmoniums with an application to information retrieval," *Proceedings of the Advances in Neural Information Processing Systems*, January 2005, Vancouver, British Columbia, Canada.
- [51] I. Sutskever and G. Hinton, "Learning multilevel distributed representations for high-dimensional sequences," *Proceedings of the International Conference on Artificial Intelligence and Statistics*, vol. 2, March 2007, San Juan, PR.
- [52] S. Ravanbakhsh, B. Póczos, J. Schneider, D. Schuurmans, and R. Greiner, "Stochastic neural networks with monotonic activation functions," *Proceedings of the International Conference on Artificial Intelligence and Statistics*, May 2016, Cadiz, Spain.
- [53] Y.-M. Huang, A. Panahi, H. Krim, Y. Yu, and S. L. Smith, "Deep adversarial belief networks," *arXiv:1909.06134*, September 2019.
- [54] Sentinel Program, "Sentinel missions," Accessed 2020.
- [55] S. Ogawa and H. Mori, "A gaussian-gaussian-restricted-boltzmann-machine-based deep neural network technique for photovoltaic system generation forecasting," *Proceedings of the International Federation of Automatic Control Workshop on Control of Smart Grid and Renewable Energy Systems*, pp. 87–92, June 2019, Jeju, Republic of Korea.
- [56] J. Hérault and B. Añs, "Neuronal network with modifiable synapses: decoding of composite sensory messages under unsupervised and permanent learning," *Comptes rendus de l'Académie des sciences. Série III, Sciences de la vie*, vol. 299, pp. 525–8, February 1984.
- [57] M. O. R. Karakida and S. Amari, "Dynamical analysis of contrastive divergence learning," *Neural Network*, vol. 79, no. C, pp. 78–87, July 2016.
- [58] K. Pearson, "On lines and planes of closest fit to systems of points in space," *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 2, no. 11, pp. 559–572, 1901.
- [59] A. J. S. B. Scholkopf, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*, MIT Press, Cambridge, MA, USA, 2001.
- [60] J. Masci, U. Meier, D. Cireunifiedan, and J. Schmidhuber, "Stacked convolutional auto-encoders for hierarchical feature extraction," *Proceedings of the International Conference on Artificial Neural Networks*, p. 52–59, June 2011, Espoo, Finland.
- [61] L. Ruff, R. Vandermeulen, N. Goernitz, L. Deecke, S. A. Siddiqui, A. Binder, E. Müller, and M. Kloft, "Deep one-class classification," *Proceedings of the International Conference on Machine Learning*, vol. 80, July 2018, Stockholm, Sweden.
- [62] C. Zhou and R. C. Paenroth, "Anomaly detection with robust deep autoencoders," *Proceedings of the ACM SIGKDD*, p. 665–674, August 2017, New York, NY.