

Hue Modification Localization By Pair Matching

Quoc-Tin Phan
DISI, University of Trento, Italy

Michele Vascotto
DISI, University of Trento, Italy

Giulia Boato
DISI, University of Trento, Italy

Abstract—Hue modification is the adjustment of hue property on color images. Conducting hue modification on an image is trivial, and it can be abused to falsify opinions of viewers. Since shapes, edges or textural information remains unchanged after hue modification, this type of manipulation is relatively hard to be detected and localized. Based on the fact that small patches inherit the same Color Filter Array (CFA) configuration and demosaicing, any distortion made by local hue modification can be detected by patch matching within the same image. In this paper, we propose to localize hue modification by means of a Siamese neural network specifically designed for matching two inputs. By crafting the network outputs, we are able to form a heatmap which potentially highlights malicious regions. Our proposed method deals well not only with uncompressed images but also with the presence of JPEG compression, an operation usually hindering the exploitation of CFA and demosaicing artifacts. Experimental evidences corroborate the effectiveness of the proposed method.

Index Terms—Hue modification, patch matching, Siamese network

I. INTRODUCTION

Modern photography is losing its innocency due to the diversified use of image manipulation software, which allows even unexperienced users to modify digital images in different ways. Image contents are characterized mainly by geometric information like texture, edges and shapes, and by color information. Color modifications, even if does not effect geometric details, deceive human perception. They are very easy to be performed, and hard to be detected if implemented carefully.

In this paper, we address the problem of local hue modification, which is defined as the adjustment of angular position on the color circle (or color wheel) within an image area. Figure 1 illustrates hue modification by different angles ¹.



Fig. 1. Hue modification by different angles. 0° means no modification. Better viewed in color.

To cope with local image manipulations, previous works seek for artifacts of Color Filter Array (CFA) and camera sensor pattern noise. Vast camera sensors employ a CFA, where each sensor element captures the light at a certain wavelength corresponding to a color component. The remaining color components at blind positions are interpolated

from surrounding pixels. This interpolation is referred to as *demosaicing*. Image manipulations will likely generate some local or global disturbances which are inconsistent to our ordinary demosaicing artifacts [1]. More blind way to detect local disturbances is the extraction of statistical features of rich models capturing different types of neighboring dependencies [2]. These features have been proved to be effective in manipulation detection and localization, see for instance [3], [4]. Besides demosaicing, the imperfections of camera sensors also create sort of camera fingerprint, the so-called Photo-Response Nonuniformity (PRNU) noise, which is supposed to be present in every image [5]. In the presence of manipulation, this pattern noise is distorted and this distortion can be exploited as a useful clue, provided that the reference PRNU can be reliably estimated and the forged region is sufficiently large.

The specific local image manipulation considered in this paper, hue modification, distorts artifacts of demosaicing and neighboring dependencies. Based on this fact, the pioneering work in [6] analyzes demosaicing artifacts and then estimates hue modification. Essentially, an interpolated value is bigger than the minimum and smaller than the maximum of its neighborhood. We can count the number of pixel values satisfying this constraint at recorded and empty positions of CFA, and obtain a big ratio between the two quantities. The estimation of hue modification is therefore done via searching over a set of modification angles until the aforementioned ratio is maximized. We want to point out that CFA analysis requires the knowledge of CFA configuration, at least the positions of green component. Such information is not always available, especially for online images. Moreover, when the image undergoes JPEG compression, demosaicing artifacts are significantly distorted. Differently, the method proposed in [7] and [8] recovers the modification angle, by modifying the questioned image with a set of angles and matching its residual with the reference PRNU. In real scenarios this technique is very difficult to be exploited since the assumption to know the reference PRNU (or have access to images to estimate it) is very strong and cannot be easily satisfied.

In this work, we propose a novel method for detecting hue modification. Our methodology exploits the fact that two patches on the same image have the same inherent CFA configuration and demosaicing. Hue modification on a local region creates inconsistencies with the rest of the image, and thus pair-wise patch matching can reveal the forged region. To achieve such purpose, we propose a solution based on Siamese neural networks [9], trained on positive pairs (two pristine patches) and negative pairs (a pristine and a modified

¹https://www.imagemagick.org/Usage/color_mods/, last access: 15/02/2019

patch). JPEG compression before and after hue modification is included during training, granting the network the capability to deal with real-world conditions. Finally, we fuse multiple outputs of patch matching to obtain a unique decision map (heatmap), on which a postprocessing is applied to precisely localize the forged region (Section II). Experiments demonstrate the effectiveness of the proposed approach (Section III).

II. PROPOSED METHOD

Hue modification is performed in the HSV space by adding an angle α to the value of H . Since defined on a circle, hue modification is periodic with a period 360° , i.e. a modification of α is identical to $\alpha \pm 360^\circ$. Besides hue, other attributes of a color in HSV space are saturation and value (brightness), whose changes are different from hue modification. Here, we investigate the detection of hue modification on: i) uncompressed images, ii) JPEG images where the modification is carried out before and after compression.

Given two rectangular patches $\mathcal{P}_i, \mathcal{P}_j$ of size $h \times w$ from the same image, we desire to estimate the logistic prediction p_{ij} that two patches are inconsistent with respect to two corresponding modification angles α_i, α_j . The two patches are *consistent* if $\alpha_i = \alpha_j = 0$ and *inconsistent* if $\alpha_i \neq \alpha_j = 0$.

We propose to verify the inconsistency of \mathcal{P}_i and \mathcal{P}_j by means of a Siamese neural network [9]. Siamese neural networks have been recently exploited for applications in multimedia forensics [10]–[12]. This network architecture consists of two identical sub-networks f_θ , followed by a non-linear classifier g_γ that outputs an inconsistency score z_{ij} whose standard logistic activation is defined as:

$$p_{ij} = \frac{1}{1 + e^{-z_{ij}}}, \text{ and } z_{ij} = g_\gamma \left([f_\theta(\mathcal{P}_i) - f_\theta(\mathcal{P}_j)]_{\text{pointwise}}^2 \right).$$

The network parameters θ, γ are jointly optimized to minimize the binary cross-entropy of network logistic predictions p_{ij} and patch inconsistencies y_{ij} , written in terms of a loss function over N training patches:

$$\mathcal{L} = \frac{1}{N^2 - N} \sum_{i=1}^N \sum_{j=1, j \neq i}^N -y_{ij} \log(p_{ij}) - (1 - y_{ij}) \log(1 - p_{ij}).$$

In Figure 2, we provide the sketch of the network architecture. We use the 50-layer Residual Network (ResNet50) [13] as the feature extractor f_θ , which outputs a 256-dimensional feature vector. The inconsistency of features extracted from two patches are evaluated by a pointwise squared difference operator. The classifier g_γ is a multilayer perceptron network composed by one hidden layer of 16 units and one single-unit output layer with sigmoid activation outputting p_{ij} . We train two separate Siamese networks end-to-end on large-scale synthetic training sets. The first model is trained on 400,000 64×64 patches extracted from uncompressed images from RAISE [14] and Dresden [15]. To train the second model, we use the same 400,000 patches and perform hue modification before or after JPEG compression with random quality factors in [55, 100]. Parameters θ are initialized using ResNet50 pretrained on ImageNet [16]. On each training

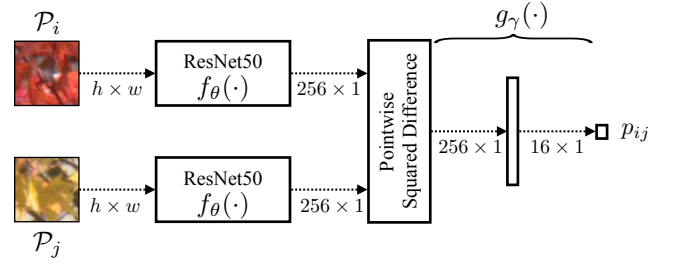


Fig. 2. Proposed Siamese network architecture.

iteration, we optimize the loss function with respect to (θ, γ) on a mini-batch of 64 pairs, half of which is labeled as positive, i.e. both two patches are unmodified, and another half is labeled as negative, i.e. one patch is modified by an angle randomly selected in $[30, 330]$ with step 30, and its counterpart is unmodified. Hue modification and JPEG compression are carried out during training. We use Adam optimizer with the starting learning rate 10^{-4} , and schedule to halve it every 5 epochs after the first 30 epochs until convergence.

A. Detection and Localization

1) *Heatmap creation:* The described architecture outputs the logistic patches inconsistency. Given a test image, we collect all inconsistency scores and generate a unique localization heatmap which potentially indicates malicious regions.

Let H, W be height and width of the image, and h, w be height and width of the small patch. By using a sliding window with stride s , the total number of patches will be $N = N_H \times N_W$, where $N_H = \lfloor \frac{H-h}{s} \rfloor + 1$ and $N_W = \lfloor \frac{W-w}{s} \rfloor + 1$ are number of patches along each dimension.

Generally, computing inconsistency scores on all possible pairs is expensive because the number of pairs grows quadratically w.r.t. N . Nevertheless, almost computational burden is attributed to operations of feature extraction network f_θ which composes convolutional layers. In pairwise manner, one patch is paired with other $N - 1$ patches and passed through f_θ about $N - 1$ times. This redundancy can be reduced by first pre-extracting low-dimensional features of N patches by evaluating $f_\theta(\mathcal{P}_i)$, $1 \leq i \leq N$, and afterwards compute p_{ij} for all possible pairs using computed features.

For each patch \mathcal{P}_k within the image, an inconsistency map $\mathcal{I}^k \in \mathbb{R}^{N_H \times N_W}$ is built. If we consider all patches according to their spatial location on the image, \mathcal{I}_{ij}^k is the inconsistency of (i, j) -th patch and \mathcal{P}_k , where $1 \leq i \leq N_H$ and $1 \leq j \leq N_W$.

It is typical to assume that the forged region is relatively small compared to the background, thus majority of \mathcal{I}^k (k refers to patches on the pristine region) exposes inconsistencies with the forged region, while remaining maps expose inconsistencies with the pristine region, as shown in Figure 3. In order to fuse inconsistency maps of majority patches belonging to the pristine region to obtain a unique map $\bar{\mathcal{I}} \in \mathbb{R}^{N_H \times N_W}$, we follow the approach in [12], computing $\bar{\mathcal{I}}$ by mean shift algorithm [17], which iteratively finds mean of majority (mode). Eventually, $\bar{\mathcal{I}}$ is a subsampled heatmap which potentially highlights malicious region. The full-size heatmap




Fig. 3. Fusing patch-level inconsistency maps into an image-level map.

can be obtained by resizing $\bar{\mathcal{I}}$ with bilinear interpolation. If the forged region is larger than the background, we obtain the inverted heatmap since the background is the smaller area.

2) *Postprocessing*: The standard logistic output p_{ij} can be interpreted as the posterior probability that two patches \mathcal{P}_i and \mathcal{P}_j are inconsistent. After mean shifting, each element $\bar{p}_{ij} = \bar{\mathcal{I}}_{ij}$ tells us how probable (i, j) -th patch is forged because $\bar{\mathcal{I}}$ is the representative inconsistency map of pristine patches to all patches. While the threshold 0.5 may be a reasonable choice for deciding if two patches are inconsistent, it is not straightforward to apply this rule to pixel-level predictions. Moreover, as keeping False Alarm Rate (FAR) low is critical in forensic applications, a postprocessing step is important for pixel-level prediction. With this respect, postprocessing on each image is cast to finding a statistical threshold τ based on which a pixel is masked as forged or pristine. We apply a simple postprocessing based on the assumption that $\bar{\mathcal{I}}_{ij}$ (to avoid adding new notation, we mean $\bar{\mathcal{I}}$ after resized) follows a Gaussian distribution, $\bar{\mathcal{I}}_{ij} \sim N(\mu, \sigma^2)$. We fix τ such that 5% of the right tail are decided as being forged. τ is lower bounded by 0.5 to maintain acceptable FAR, namely $\tau = \max(0.5, t)$. t is the solution of: $0.95 = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^t e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx$. Compared to the threshold 0.5, τ results in better or equal FAR. An example of postprocessing is shown in Figure 4.




Fig. 4. The heatmap before (left) and after (right) being postprocessed. The middle plot presents CDF of $\bar{\mathcal{I}}_{ij}$ and the threshold $\tau = 0.61$.

III. EXPERIMENTS

Towards experimental evidences, we evaluate our approach under different configurations and test sets.

A. Test set

To the best of our knowledge, there is no publicly available dataset on the problem of hue modification. Thus, to evaluate our method, we generated the test set from 120 raw images of an external Canon 600D camera (never appeared in training phase) having CFA pattern *GBRG*. Raw images are decoded by *dcraw* version 9.27. For each image, a top-left region is cropped out such that $H = 768$ and $W = 1024$. The forged

area follows random convex shape fixed within a 256×256 bounding box, which is positioned at random location on the image. Next, we perform hue modification on pixels inside the polygons and generate multiple test sets:

- \mathcal{D}_{png}^α : Uncompressed images are demosaiced from raw images by *dcraw* and subject to local hue modification. For each modification angle $\alpha \in [30, 330]$ step 30, hue modification is carried out on 120 uncompressed images.
- \mathcal{D}_{b-jpg}^{QF} : Hue modification by different angles (10 images for each modification angle $\alpha \in [30, 330]$ step 30) are carried out on 110 images, and 10 images are unmodified. Afterwards, all images are compressed using quality factors $QF \in [55, 100]$, step 5.
- \mathcal{D}_{a-jpg}^{QF} : 120 images are first compressed using $QF \in [55, 100]$, step 5. Afterwards, hue modification by different angles (one angle for 10 images) are carried out on 110 JPEG images, while the remaining 10 images are unmodified. All of them are compressed again using the default quality factor 75. By the second JPEG compression, \mathcal{D}_{a-jpg}^{QF} is more challenging since the training images are only subject to single JPEG compression.

B. Setups

The performance of our method is compared with the following state of the art methods: Choi *et al.* [6], based on CFA-based artifacts and explicitly designed for the estimation of hue modification, and SpliceBuster [3], based on statistical features of rich models [2] and selected for comparison since those features potentially capture local disturbances caused by local hue modification. We do not compare with [7], [8] given their strong assumption about the availability of the reference PRNU which is unrealistic in practical scenarios.

This work particularly focuses on the localization of hue modification rather than its estimation. Choi *et al.* [6] is an estimator which potentially returns the modification angle by searching over a feasible range. To convert Choi *et al.* into a localization method, we use a sliding window 35 similarly to our method, and search the angle over $[0, 359]$, step 8. If the angle found is 0 or 352, the patch is marked as pristine. Choi *et al.* therefore outputs a binary map. The other method, SpliceBuster [3], returns the negative log-likelihood that a pixel is pristine. It means, a large value indicates high probability that a pixel is forged. We linearly scale the returned map into $[0, 1]$ and apply the same postprocessing described in Section II-A2 to get the binary map. In order to demonstrate the advantage of our postprocessing, we also report performance of the proposed method when a simple thresholding is applied to binarize the heatmap. We empirically found that the threshold 0.8 yields most acceptable results.

We aggregate True Positive (TP), True Negative (TN), False Positive (FP), and False Negative (FN) over all images and report average True Positive Rate (TPR), True Negative Rate (TNR) and F1 score.

TABLE I
F1 SCORES OF ALL METHODS ON \mathcal{D}_{png}^α .

Angle $\alpha \rightarrow$ Method \downarrow	30	90	150	210	270	330
Choi <i>et al.</i>	66.16	67.15	65.05	68.34	67.92	68.21
SpliceBuster	18.12	29.57	32.99	26.80	21.97	12.29
Siamese-T-0.8	66.44	65.28	66.22	69.28	70.00	63.23
Siamese-G-0.95	71.79	73.82	73.04	74.41	74.29	69.86

C. Quantitative Evaluation

1) *Detection on uncompressed images:* In this section, we evaluate the first model trained on uncompressed images, and compare with Choi *et al.* and SpliceBuster on \mathcal{D}_{png}^α . Figure 5 presents TPR and TNR obtained by all methods. Choi *et*




Fig. 5. TPR and TNR of all methods on \mathcal{D}_{png}^α .

al. is guaranteed to detect hue modification on uncompressed images since this type of manipulation distorts demosaicing artifacts. It achieves high TPR which implies that almost forged pixels are correctly detected. This comes at a cost of slightly worse TNR. SpliceBuster, on the other hand, detects correctly only about 20% of forged pixels, and as a consequence, yields very high TNR. We assume that the features used in [2] are ineffective for hue modification detection.

Our Siamese network with heatmaps thresholded simply by 0.8 is denoted by Siamese-T-0.8. The other alternative is denoted by Siamese-G-0.95, where heatmaps are postprocessed by threshold τ as designed in Section II-A2. We can clearly see that Siamese-G-0.95 outperforms Siamese-T-0.8 in all cases. In fact, a fixed threshold over all heatmaps cannot deal with high variability of predictive scores on each heatmap, and thus an adaptive threshold is more effective. Interestingly, the TPR reveals the fact that the middle range of modification angles are easier to detect by our methodologies. This is explainable since the strength of hue modification is periodic with the period of 360° . Very small or very large positive angles correspond to little modifications.

We summarize the overall performance for some selective modification angles in Table I. In terms of F1 score, Siamese-G-0.95 outperforms all other methods.

2) *Detection in the presence of JPEG compression:* We target more practical scenarios where hue modification is done with the presence of JPEG compression. It has been acknowledged that JPEG compression has strong impact on demosaicing artifacts [18]–[20]. Choi *et al.* is also very sensitive to JPEG compression since the count of interpolated and recorded pixels is less accurate [6].

TABLE II
F1 SCORES OF ALL METHODS ON \mathcal{D}_{b-jpg}^{QF} .

$QF \rightarrow$ Method \downarrow	75	80	85	90	95	100
Choi <i>et al.</i>	9.80	9.74	9.69	9.85	10.05	39.47
SpliceBuster	12.80	12.47	14.26	14.66	15.57	21.44
Siamese-T-0.8	51.96	53.43	54.77	54.56	56.81	61.60
Siamese-G-0.95	61.41	60.83	63.03	64.65	66.73	69.11

We assess the second model trained on JPEG images under two testing circumstances: i) hue modification is performed on uncompressed images followed by JPEG compression, and ii) hue modification is performed on JPEG compressed images, and those are subsequently compressed again using quality factor 75. Note that during training, we do not perform second JPEG compression.

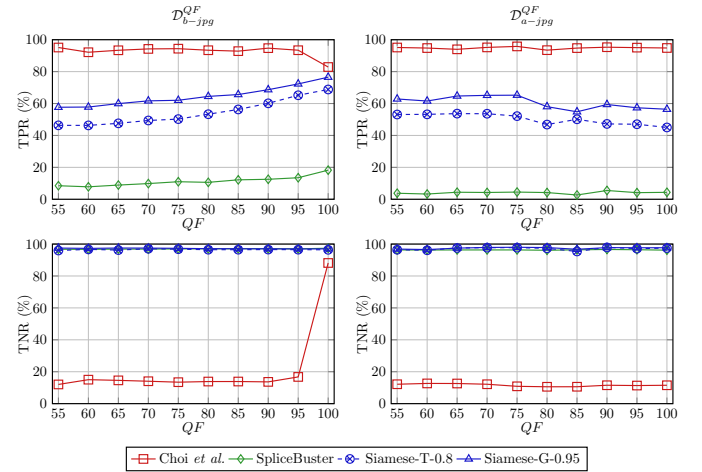


Fig. 6. TPR and TNR of all methods on \mathcal{D}_{b-jpg}^{QF} and \mathcal{D}_{a-jpg}^{QF} . QF is the first compression quality factor.

The TPR and TNR of all methods, where hue modification is performed before JPEG compression, i.e. dataset \mathcal{D}_{b-jpg}^{QF} , is shown in the first column of Figure 6. Choi *et al.* fails to spot forged area unless the image is compressed with highest QF . At $QF = 100$, Choi *et al.* achieves TPR 82.88% and TNR 88.15%. SpliceBuster, on the other hand, can only detect about 10% of forged pixels. Our proposed methods perform far better than the other two competitors on \mathcal{D}_{b-jpg}^{QF} , by keeping TPR at acceptable level and retaining always high TNR, i.e., Siamese-G-0.95 achieves an average 64.67% of TPR ($\approx 65\%$ of forged pixels are correctly detected) and 97.72% of TNR, while Siamese-T-0.8 attains 54.35% of TPR and 97.39% of TNR. In the right column of Figure 6, i.e. dataset \mathcal{D}_{a-jpg}^{QF} , the overall TPR and TNR of our methods are slightly degraded compared to the performance on \mathcal{D}_{b-jpg}^{QF} . This degradation can be attributed to the second JPEG compression. In fact, we can generally observe the correlation of performance degradation and compression rate: the higher the first QF , the lower the performance. While Choi *et al.* behaves positively on $\mathcal{D}_{b-jpg}^{100}$, it loses that capability on $\mathcal{D}_{a-jpg}^{100}$.

The overall F1 scores for several selective QF are shown in

TABLE III
F1 SCORES OF ALL METHODS ON \mathcal{D}_{a-jpg}^{QF} .

$QF \rightarrow$ Method \downarrow	75	80	85	90	95	100
Choi <i>et al.</i>	9.68	9.62	9.55	9.70	9.65	9.65
SpliceBuster	5.07	4.63	3.08	6.35	4.74	4.82
Siamese-T-0.8	52.85	46.76	41.29	49.93	46.13	46.09
Siamese-G-0.95	63.51	57.52	49.83	58.36	56.32	55.75

Table II and III. Our two methodologies, in particular Siamese-G-0.95, outperform the other two methods by a large margin. We might notice that Choi *et al.* achieves F1 score 39.47% on $\mathcal{D}_{b-jpg}^{100}$ while TPR and TNR in the same configuration are over 80%, see left column in Figure 6. This phenomenon is due to the high FP which penalizes precision, and as a consequence, F1 score. However, since TN dominates FP (due to the large pristine area compared to the forged area), TNR is not effectively penalized.

D. Qualitative Inspection

In Figure 7, we provide detection results on realistic examples manually created using GIMP. Hue modification is carried out on uncompressed images (the first 5 lines) and the modified images are JPEG compressed using highest quality $QF = 100$ (the last 5 lines). Siamese-G-0.95 (the last column) clearly results in better detection maps compared with Choi *et al.* and SpliceBuster.

IV. CONCLUSION

We have proposed a data-driven countermeasure for hue modification on color images based on patch matching. This task is done by means of a Siamese architecture which receives the two inputs and outputs the likelihood that the two inputs are inconsistent. A unique localization map is generated from inconsistency scores of multiple patches. Our models perform well on uncompressed and JPEG compressed images even though JPEG compression distorts CFA and demosaicing artifacts. Our future investigations will focus on the estimation of hue modification angles, based on which the original image can be recovered.

REFERENCES

- [1] P. Ferrara, T. Bianchi, A. De Rosa, and A. Piva, "Image forgery localization via fine-grained analysis of CFA artifacts," *IEEE Trans. on Information Forensics and Security*, vol. 7, no. 5, pp. 1566–1577, 2012.
- [2] J. Fridrich and J. Kodovsky, "Rich models for steganalysis of digital images," *IEEE Trans. on Information Forensics and Security*, vol. 7, no. 3, pp. 868–882, 2012.
- [3] D. Cozzolino, G. Poggi, and L. Verdoliva, "Splicebuster: A new blind image splicing detector," in *Proc. of WIFS*, 2015, pp. 1–6.
- [4] H. Li, W. Luo, X. Qiu, and J. Huang, "Image forgery localization via integrating tampering possibility maps," *IEEE Trans. on Information Forensics and Security*, vol. 12, no. 5, pp. 1240–1252, 2017.
- [5] M. Chen, J. Fridrich, M. Goljan, and J. Lukáš, "Determining image origin and integrity using sensor noise," *IEEE Trans. on Information Forensics and Security*, vol. 3, no. 1, pp. 74–90, 2008.
- [6] C.-H. Choi, H.-Y. Lee, and H.-K. Lee, "Estimation of color modification in digital images by CFA pattern change," *Forensic science international*, vol. 226, pp. 94–105, 01 2013.
- [7] J. Hou, H. Jang, and H. Lee, "Hue modification estimation using sensor pattern noise," in *Proc. of ICIP*, 2014, pp. 5287–5291.




Fig. 7. Detection results on realistic examples.

- [8] J. Hou and H. Lee, "Detection of modification using photo response nonuniformity," *IEEE Trans. on Circuits and Systems for Video Technology*, vol. 27, no. 8, pp. 1826–1832, 2017.
- [9] J. Bromley, I. Guyon, Y. LeCun, E. Säckinger, and R. Shah, "Signature verification using a "siamese" time delay neural network," in *Proc. of NIPS*, 1993, pp. 737–744.
- [10] D. Cozzolino and L. Verdoliva, "Noiseprint: a CNN-based camera model fingerprint," *IEEE Trans. on Information Forensics and Security (Early Access)*, 2019.
- [11] O. Mayer and M. C. Stamm, "Learned forensic source similarity for unknown camera models," in *Proc. of ICASSP*, 2018, pp. 2012–2016.
- [12] M. Huh, A. Liu, A. Owens, and A. A. Efros, "Fighting fake news: Image splice detection via learned self-consistency," in *Proc. of ECCV*, 2018.
- [13] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. of CVPR*, 2016, pp. 770–778.
- [14] D.-T. Dang-Nguyen, C. Pasquini, V. Conotter, and G. Boato, "RAISE: A raw images dataset for digital image forensics," in *Proc. of ACM MMSys*, 2015, pp. 219–224.
- [15] T. Gloe and R. Böhme, "The 'Dresden Image Database' for benchmarking digital image forensics," in *Proc. of ACM SAC*, 2010, vol. 2, pp. 1585–1591.
- [16] J. Deng, W. Dong, R. Socher, L. Li, and and, "ImageNet: A large-scale hierarchical image database," in *Proc. of CVPR*, 2009, pp. 248–255.
- [17] Y. Cheng, "Mean shift, mode seeking, and clustering," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 17, no. 8, pp. 790–799, 1995.
- [18] A. C. Gallagher and T.-H. Chen, "Image authentication by detecting traces of demosaicing," in *Proc. of IEEE Workitorial on Vision of the Unseen (in conjunction with CVPR)*, 2008, pp. 1–8.
- [19] H. Cao and A. C. Kot, "Accurate detection of demosaicing regularity for digital image forensics," *IEEE Trans. on Information Forensics and Security*, vol. 4, no. 4, pp. 899–910, 2009.
- [20] M. Kirchner, "Efficient estimation of CFA pattern configuration in digital camera images," in *Proc. of SPIE, Media Forensics and Security II*, 2010, vol. 7541.