

A Method for Identifying Origin of Digital Images Using a Convolutional Neural Network

Rong Huang^{1,2}, Fuming Fang³, Huy H. Nguyen⁴, Junichi Yamagishi^{3,4}, Isao Echizen^{3,4,5}

¹ College of Information Science and Technology, Donghua University, Shanghai, China

² Engineering Research Center of Digitized Textile & Apparel Technology, Ministry of Education, Donghua University, Shanghai, China

³ National Institute of Informatics, Tokyo, Japan

⁴ SOKENDAI (The Graduate University for Advanced Studies), Kanagawa, Japan

⁵ The University of Tokyo, Tokyo, Japan

Abstract—The rapid development of deep learning techniques has created new challenges in identifying the origin of digital images because generative adversarial networks and variational autoencoders can create plausible digital images whose contents are not present in natural scenes. In this paper, we consider the origin that can be broken down into three categories: natural photographic image (NPI), computer generated graphic (CGG), and deep network generated image (DGI). A method is presented for effectively identifying the origin of digital images that is based on a convolutional neural network (CNN) and uses a local-to-global framework to reduce training complexity. By feeding labeled data, the CNN is trained to predict the origin of local patches cropped from an image. The origin of the full-size image is then determined by majority voting. Unlike previous forensic methods, the CNN takes the raw pixels as input without the aid of “residual map”. Experimental results revealed that not only the high-frequency components but also the middle-frequency ones contribute to origin identification. The proposed method achieved up to 95.21% identification accuracy and behaved robustly against several common post-processing operations including JPEG compression, scaling, geometric transformation, and contrast stretching. The quantitative results demonstrate that the proposed method is more effective than handcrafted feature-based methods.

I. INTRODUCTION

Nowadays, with the proliferation of powerful computer graphics tools, such as Softimage XSI, Maya, and TerraGen, it has become easy for non-professional technicians to artificially create digital scenes without leaving any perceptible clues. Figures 1 (a) and (b) show a natural photographic image (NPI) and a computer generated graphic (CGG), respectively. At first glance, it is not obvious that the second image was generated by computers. In media forensics, substantial efforts have been made to distinguish between a CGG and an NPI. Statistical clues based on wavelet coefficients [1], histograms [2], edge pixels [3], texture [4], entropy [5], and multi-fractal spectrum [6] have been used to design handcrafted features geared towards a particular classifier. Furthermore, pixel-level inconsistencies such as JPEG compression artifacts [7], demosaicking clues [8], lens aberration [9], and sensor pattern noise [10] are useful indicators for the task of identifying origin. Recently, deep neural networks have been the workhorse for a wide variety of computer vision tasks, including image classification, image annotation, and

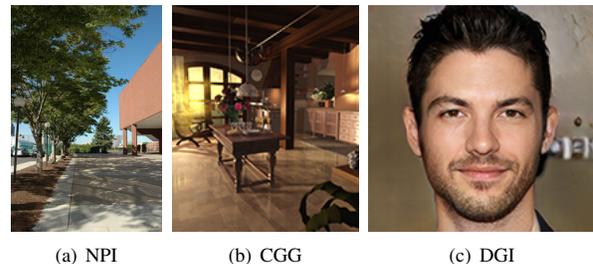


Fig. 1. Example digital images with different origins.

object detection. Driven by data, CNNs have the ability to automatically learn hierarchical representations, and thereby achieve better generalization in an end-to-end manner. Hence, CNN-based methods [11–13] have been quickly adopted by the multimedia security community for distinguishing CGGs from NPIs.

Deep neural networks, however, are like a double-edged sword: they can be used not only for identifying origin but also for generating photorealistic images if the network is structured as a generative adversarial network (GAN) or variational autoencoder (VAE). Since most generative models can be trained in an unsupervised manner, it is not prohibitively expensive for an attacker to generate fake scenes that can be maliciously used for illegal purposes. As evidenced by the deep network generated image (DGI) in Fig.1 (c), it is not trivial to identify DGIs with the naked eye. Several countermeasures have been proposed recently. McCloskey and Albright [14] found that the frequency of saturated and under-exposed pixels is suppressed by the generator’s normalization steps, which provide useful discriminative traces that enable a DGI to be distinguished from an NPI. Marra *et al.* [15] revealed that specific correlation exists in the noise residual of a DGI, which can be viewed as evidence of the origin of digital images. Li *et al.* [16] exploited the disparities in color components and designed a 588-dimensional feature vector based on multiple color co-occurrence matrices. In addition, some CNN-based methods [17–20] attempt to identify DGIs and/or localize fake regions for facial images.

According to the above survey, we conclude that digital

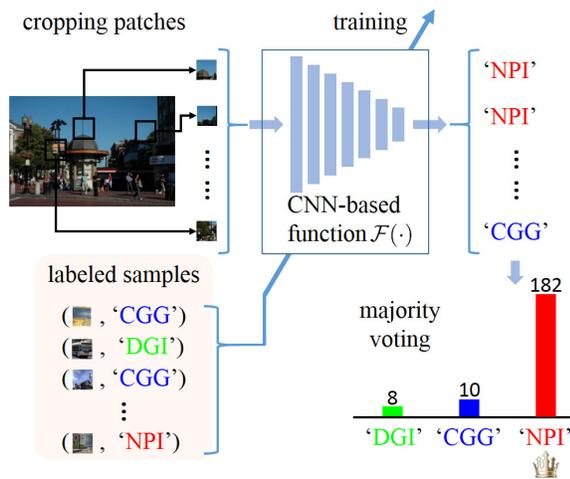


Fig. 2. Flowchart of proposed method.

images are created by one of three ways: (i) being captured with a CCD or CMOS sensor, which counts the photons passing through a color filter array, (ii) being synthesized using model-based rendering software, and (iii) being generated using data-driven deep learning model (e.g., GAN and VAE). However, the existing methods based on hand-crafted features [1–10, 14–16] can only distinguish a CGG or DGI from an NPI. Therefore, an end-to-end method that can automatically probe discriminative features and complete the origin identification task is needed to ensure the information security of digital images. To achieve this, we develop a CNN-based method in a local-to-global framework. The experimental results demonstrate that the proposed method is better in terms of identification accuracy and robustness compared with the handcrafted feature-based methods [2, 14, 16].

II. PROPOSED METHOD

The identification task considered in this paper is a three-class problem. Suppose that a training set comprises N samples along with their labels $\{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n), \dots, (x_N, y_N)\}$, where $n = 1, 2, \dots, N$. The notation x_n represents the n^{th} to-be-identified sample while y_n corresponds to its label. Our goal is to exploit the training set and establish a CNN-based mapping function $\mathcal{F}(\cdot)$ so as to make correct predictions for new samples.

We use a local-to-global framework to reduce training complexity. Specifically, for an RGB color image, we crop M local patches, each of size $224 \times 224 \times 3$. After the CNN model is trained, it is expected to predict the correct label, denoted by \hat{y}_n , for each local patch. The origin of the full-size image is then determined by majority voting. Figure 2 schematically illustrates the flowchart of the proposed method. Using the local-to-global framework not only reduces training complexity but also encourages the CNN model to learn forensic clues rather than to recognize the image contents.

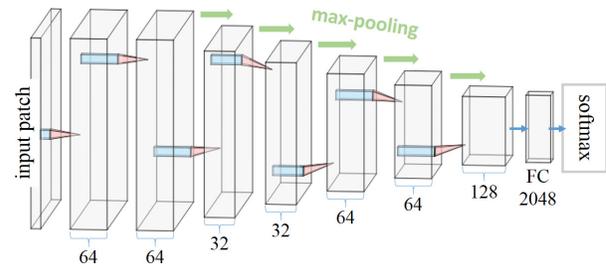


Fig. 3. Proposed CNN architecture.

A. Patch Cropping

Sampling representative data is an important prerequisite for training a statistical learning model. Motivated by the observation that highly textured areas usually contain clues about factitiousness, we preferentially crop patches containing rich edge pixels. Specifically, the edge pixels are detected by Canny detector. The high and low threshold values are set to 50 and 100, respectively. The image is then partitioned into overlapping patches with an appropriate stride to ensure that the number of candidate patches exceeds $20M$. The number of edge pixels within a candidate patch is considered to be the fitness value. From the $20M$ candidate patches, roulette wheel selection, which is widely used in genetic algorithms, is performed to choose M ones without replacement. Compared with random cropping, the roulette-based strategy can avoid cropping too many smooth patches.

B. Data Augmentation

As a regularization technique, data augmentation protects well against overfitting. The training set is artificially expanded by generating new instances from existing ones. To this end, the patches are rotated by 90° , 180° , and 270° , and horizontal and vertical flipping are performed. A total of $6M$ patches can be harvested from a full-size image by combining these geometric transformations.

Furthermore, these newly generated patches help the model to be more robust against rotation and flipping. The experimental results presented in Section III support this claim.

C. Network Architecture

CNNs typically consist of a cascade connection of multiple convolutional layers, pooling operations, and one or more fully connected (FC) layers. Each convolutional layer, which can be viewed as a group of filters, is trained to produce an appropriate feature representation. Each pooling operation, which computes a summary statistic of a certain layer's output, helps to make the feature representation approximately invariant to small translations. The learned hierarchical representation is fed into the FC layer, which plays a role in making correct predictions.

The proposed CNN architecture is shown in Fig. 3. The input local patches are $224 \times 224 \times 3$. The CNN model consists of seven convolutional layers, five pooling operations and one FC layer. Specifically, there are 64, 64, 32, 32, 64, 64, and

128 filters for the 7 convolutional layers, respectively. The filters are all 3×3 and have a unit stride. A rectified linear unit (ReLU) is used for activation. In each convolutional layer, batch normalization is used to mitigate the problem of internal covariate shift, making the model significantly easier to train. The outputs of the last five convolutional layers are further processed by a 2×2 max-pooling operation with stride 2. The final FC layer contains 2048 neural nodes followed by a three-way softmax function.

In previous steganalysis and forensics methods [16, 17], the input was pre-processed using fixed high-pass filters to suppress image contents and extract the so-called “residual map”. However, there is no evidence to suggest that the middle (or low) frequency components make no contribution to the forensics task. We therefore skip this pre-processing step and make all the filters trainable in the hope of adaptively probing useful information.

D. Training

We trained the CNN model using two scenarios. In scenario 1, the first two convolutional layers simply duplicate the previously trained ‘conv1_1’ and ‘conv1_2’ of the VGG19 model [21], and the corresponding weights are kept fixed during future training. In scenario 2, all layers are trained without any constraints by directly using task-dependent data. For convenience, we call these two scenarios ‘vgg’ and ‘ada’, respectively.

The cross-entropy between y_n and \hat{y}_n , where $n = 1, 2, \dots, N$, is defined as the cost function. A gradient-based stochastic optimization method called ADAM [22] is used to update the weights with a learning rate of 10^{-4} , a minibatch size of 128, and three default settings ($\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$). To prevent the network from overfitting, a powerful regularization technique called Dropout [23] is applied during the training phase to the FC layer by setting the inclusion probability to 0.2. We empirically stop the training process after ten epochs.

E. Dataset

We used a standard dataset called Columbia Photographic Images and Photorealistic Computer Graphics (PRCG) [24], which contains 800 CGGs collected from the Internet and 800 NPIs obtained from the personal collections. To prepare DGIs, we used a pre-trained progressive GAN (ProGAN) [25] to generate a sufficient number of high-quality images¹. Then, we manually screened out the 800 most plausible ones and grouped them into the third set. Unlike methods that focus on only facial images [17–20], our proposed method aims at the origin identification. Therefore, the images in the dataset encompass various scenes and contents including human faces, animals, landscapes, objects, plants, buildings, vehicles, etc., for each class (NPI, CGG, and DGI). This enables us to train the model and thus make it more suitable for real-world applications. We show some

¹The ProGAN was trained on the LSUN dataset [26] (category-wise) and CelebA-HQ dataset [27].

TABLE I
IDENTIFICATION ACCURACY AND PERFORMANCE COMPARISON.

| Methods | Accuracy (%) | | |
|--------------------------------|----------------|-------------|-------|
| | patch level | image level | |
| Histogram-based [2] | / | 87.50 | |
| Saturation-based [14] | | 56.67 | |
| Color co-occurrence-based [16] | | 91.67 | |
| Proposed | ‘vgg’ scenario | 90.32 | 94.38 |
| | ‘ada’ scenario | 92.90 | 95.21 |

samples of the dataset online <https://nii-yamagishilab.github.io/Samples-Rong/Attack-Type-detection/>.

III. EVALUATION

We experimentally evaluated our proposed method: (i) to determine its identification accuracy, (ii) to study which frequency components are more useful for origin identification, and (iii) to determine its robustness against several post-processing operations including JPEG compression, scaling, geometric transformation, and contrast stretching. For comparative studies, we used three handcrafted feature-based methods [2, 14, 16] as baselines.

A. Evaluation Method

In data preparation, we randomly divided the image dataset into three subsets: training (70%), validation (10%), and test (20%). From each training image, we cropped $M = 200$ local patches each of size $224 \times 224 \times 3$ and then performed geometric transformations to expand the training set. This produced $N = 2.016 \times 10^6$ training samples. The CNN model was implemented using the TensorFlow framework. The computing platform consists of an i7-6770K 4GHz CPU and a NVIDIA GeForce GTX 1070 GPU with 16GB memory.

For the three baselines, we extracted the tailor-made features and uniformly used a discriminant analysis classifier. For the histogram-based method [2], the number of bins was set to seven, as recommended by the authors. We simply calculated the average accuracy over three color channels. For the saturation-based method [14], each image was first converted from RGB into grayscale, and then an 8-dimensional descriptor was constructed by combining both the under- and over-exposed pixel frequencies. For the color co-occurrence-based method [16], the order of the co-occurrence matrix was set to 3 under a sample-aware scenario.

B. Identification Accuracy Results

Figure 4 shows three examples of origin identifications at the patch level. The outputs of the CNN model are indicated by the red, green, and blue bounding boxes, which correspond to ‘NPI’, ‘DGI’ and ‘CGG’, respectively. The panel above each resulting image shows the voting table, which contains the number of bounding boxes by color. As we saw, the majority voting stage has the ability to suppress the objectionable

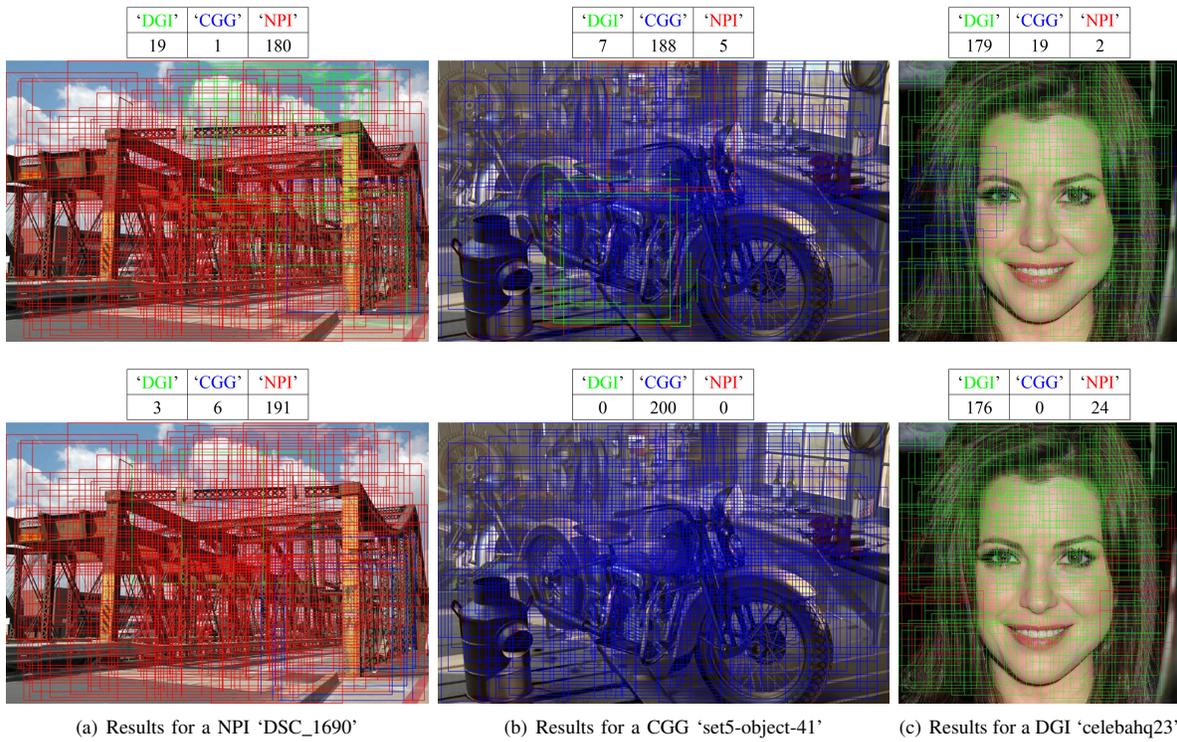


Fig. 4. Resulting examples of origin identifications at patch level and quantitative results of majority voting. First row shows results for 'vgg' scenario; second row shows results for 'ada' scenario.

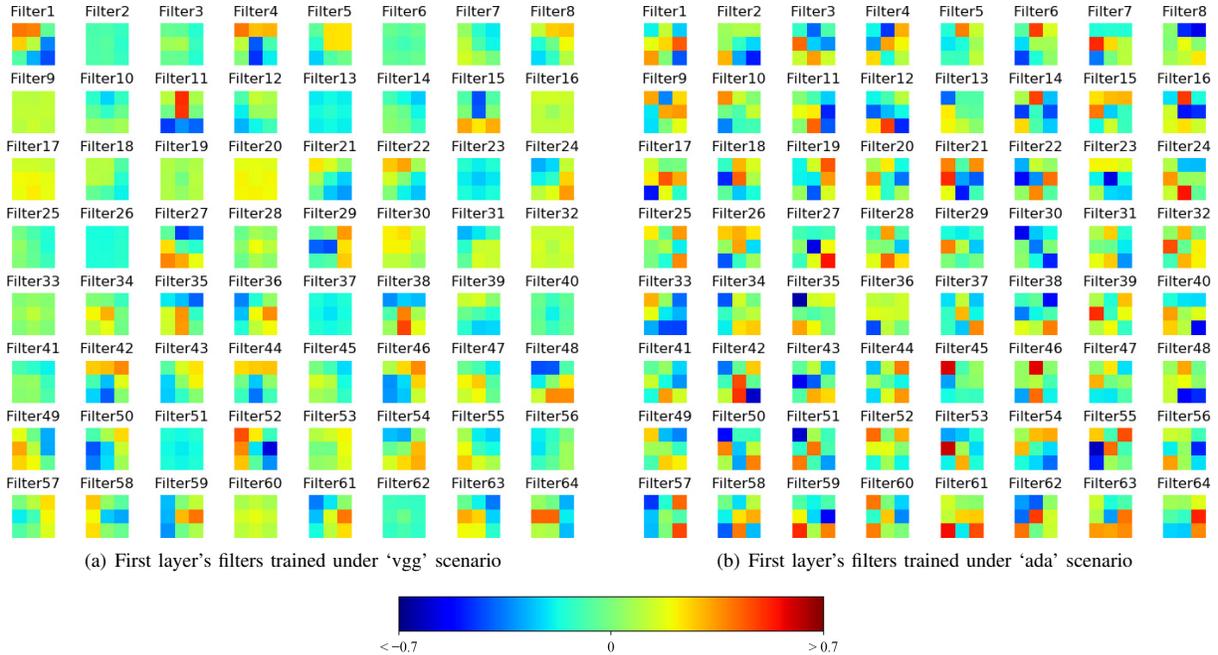


Fig. 5. Visualizing trained filters at the first convolutional layer (the third channel).

outliers, thereby improving identification performance at the image level.

The identification accuracy results obtained using the test set are summarized in Table I. Patch-level accuracy is defined

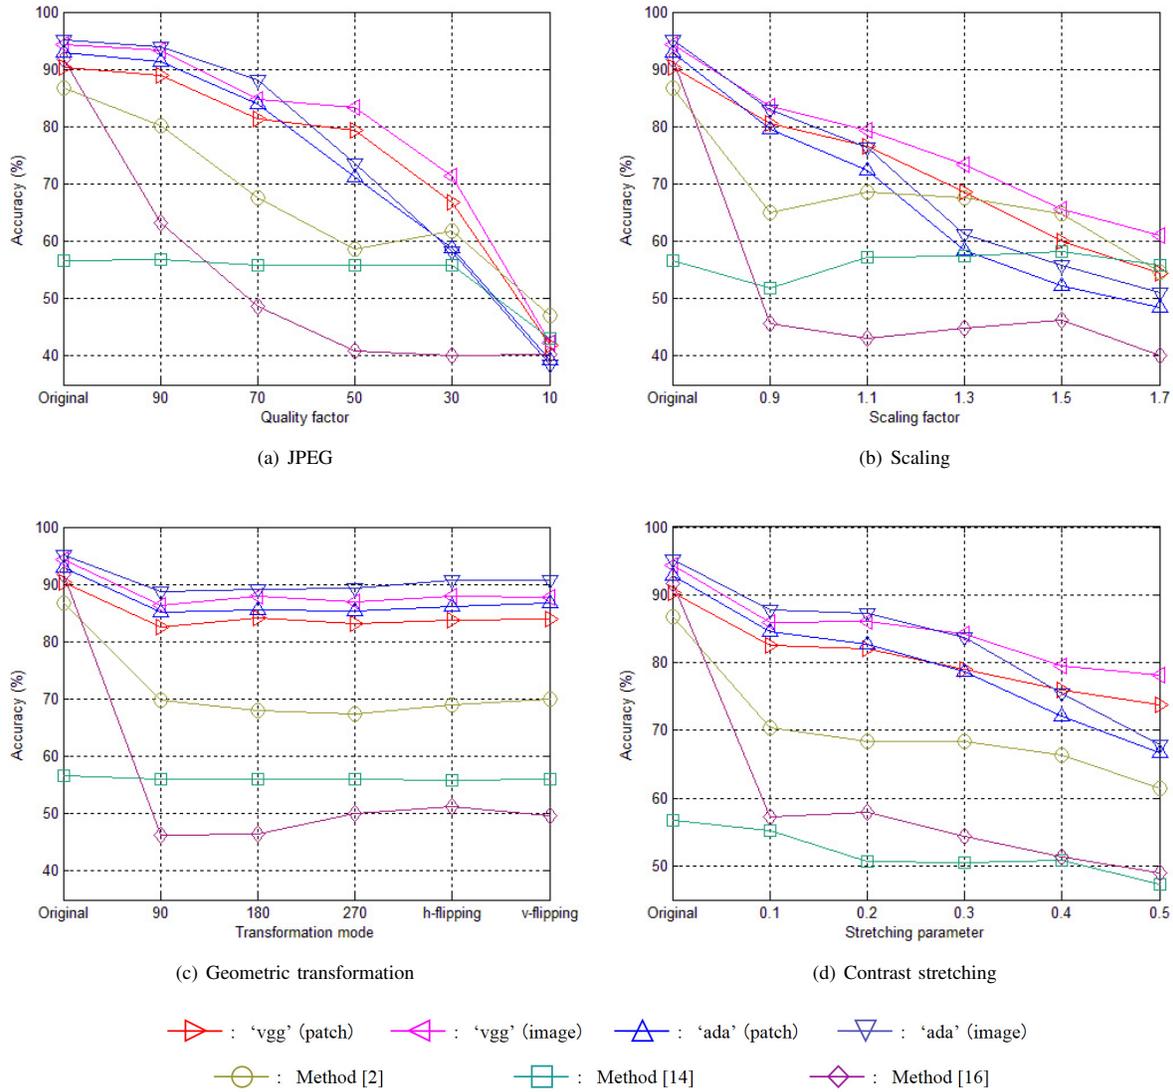


Fig. 6. Accuracy evolutions for testing robustness against JPEG compression, scaling, geometric transformation, and contrast stretching. Notation ‘h-flipping’ and ‘v-flipping’ mean horizontal and vertical flipping, respectively.

as the ratio of the number of correct predictions to the total number of test patches while image-level accuracy is defined as the ratio of the number of correct identifications to the total number of test images. The proposed CNN-based method identified image’s origin with 95.21% accuracy for the ‘ada’ scenario. The results in Table I show that our method has higher accuracy scores than the three baseline methods. In addition, the majority voting indeed boosted the identification accuracy by 4.06% for the ‘vgg’ scenario and 2.31% for the ‘ada’ scenario.

C. Frequency Components Results

To study which frequency components contain information about the origin forensic, we focused on the first layer’s filters,

which had been trained under the ‘vgg’ or ‘ada’ scenario. Figure 5 visualizes the 64 trained filters at the first convolutional layer (the third channel). We found that the ‘ada’ scenario forced the filter weights to diversify, meaning that more high-frequency components were extracted and that more image contents were suppressed. In contrast, the ‘vgg’ filters tended to extract richer frequency components. This phenomenon is also supported, in Table II, by the channel-wise variances of the filter weights in the sense that the variances for the ‘ada’ scenario were about twice those for the ‘vgg’ scenario.

As shown by the accuracy scores in Table I, the ‘ada’ scenario outperformed the ‘vgg’ one, indicating that high-frequency components possess more clues for identifying the origin. We also found that several ‘ada’ filters, like the

TABLE II
CHANNEL-WISE WEIGHT VARIANCES OF FIRST LAYER.

| | Variances | | |
|----------------|-----------|-----------|-----------|
| | channel 1 | channel 2 | channel 3 |
| 'vgg' scenario | 0.02215 | 0.03597 | 0.02999 |
| 'ada' scenario | 0.05509 | 0.06180 | 0.06357 |

47th one, were devoted to middle-frequency components. This suggests that the middle-frequency components also contribute to the origin identification.

D. Robustness Results

Since artificially created images may be further processed using common image processing operations for special purposes, robustness against various post-processing operations is highly desirable for a good identification system. We thus examined the proposed method in terms of robustness against JPEG compression, scaling, geometric transformation, and contrast stretching. To this end, we prepared new test sets. For testing the robustness against JPEG compression, each test image was compressed for a large range of quality factors from 90 to 10 with a step size of 20. For testing the robustness against scaling, we used bicubic interpolation to resize the test images with scaling factors varying in the range [0.9,1.7]. The geometric transformations included rotation (90°, 180°, 270°), and horizontal/vertical flipping. For testing the robustness against contrast stretching, we used a piecewise-linear function that expanded the range of intensity levels so as to span the full intensity range. Two locations of points ($r_{\min}, 0$) and ($r_{\max}, 255$) control the shape of the function, where r_{\min} and r_{\max} denote the minimum and maximum intensity levels in the image, respectively. We designed a stretching parameter, α , and manipulated r_{\min} and r_{\max} : $r'_{\min} = (1 + \alpha) \cdot r_{\min}$, and $r'_{\max} = (1 - \alpha) \cdot r_{\max}$. Note that patch cropping was triggered again for each post-processed test image.

We conducted the robustness comparison with the three baseline methods. The results were plotted in Fig. 6, where 'Original' indicates the accuracy scores listed in Table I. As shown in Fig. 6 (c), the trained CNN model was particularly robust against geometric transformations. This robustness is attributed to the diversity of the training set after data augmentation. Moreover, the 'vgg' scenario achieves a comeback win over the 'ada' one when the test images are severely processed. This is because the 'vgg' filters, which are trained using a large-scale data set, have better generalization to deal with diverse samples.

For all four post-processing operations, the accuracy curves for the proposed method progressively dropped while those for the histogram-based method [2] and color co-occurrence-based method [16] dramatically deteriorated. Although the saturation-based method [14] had the most stable curves, its accuracy scores were lower than 60%. Except for several extreme cases, our method outperformed the three baseline methods even when the test images had undergone post-

processing operations. Consequently, the proposed method is better in terms of robustness than the other methods [2, 14, 16].

IV. CONCLUSION

We have presented an effective method based on a convolutional neural network for identifying the origin of digital images. Roulette wheel selection is used to crop 200 patches from an image while considering edge information. After data augmentation, the CNN is trained to predict the origin at the patch level while the origin of the full-size image is determined by majority voting. We designed two training scenarios, namely the 'vgg' and the 'ada'. The experimental results show that the 'ada' filters densely extracted high-frequency features and achieved the highest accuracy 95.21%. The 'vgg' filter, however, tended to probe various frequency components and exhibited stronger robustness. Comparison of the results with those of three handcrafted feature-based methods demonstrated that the proposed method is better in terms of identification accuracy and robustness.

ACKNOWLEDGMENT

This work was partially supported by a JST CREST Grant (JPMJCR18A6, VoicePersonae project), Japan, and by MEXT KAKENHI Grants (16H06302, 17H04687, 18H04120, 18H04112, 18KT0051), Japan.

REFERENCES

- [1] S. Lyu, H. Farid, "How realistic is photorealistic," *IEEE Trans. Signal Process.*, 53(2):845-850, 2005.
- [2] R.Y. Wu, X.L. Li, B. Yang, "Identifying computer generated graphics via histogram features," in *Proc. 18th IEEE Int. Conf. Image Process.*, pp.1933-1936, 2011.
- [3] R. Zhang, R.D. Wang, T.T. Ng, "Distinguishing photographic images and photorealistic computer graphics using visual vocabulary on local image edges," in *Proc. 10th Int. Workshop Digit. Watermarking*, pp.292-305, 2011.
- [4] X.F. Wang, Y. Liu, B.C. Xu, L. Li, J.R. Xue, "A statistical feature based approach to distinguish PRCG from photographs," *Comput. Vis. Image Underst.*, 128:84-93, 2014.
- [5] H.H. Nguyen, H.Q. Nguyen-Son, T.D. Nguyen, I. Echizen, "Discriminating between computer-generated facial images and natural ones using smoothness property and local entropy," in *Proc. 14th Int. Workshop Digit. Watermarking*, pp.39-50, 2015.
- [6] F. Peng, D.L. Zhou, M. Long, X.M. Sun, "Discrimination of natural images and computer generated graphics based on multi-fractal and regression analysis," *AEU-Int. J. Electron. Commun.*, 71:72-81, 2017.
- [7] T. Bianchi, A. Piva, "Image forgery localization via block-grained analysis of JPEG artifacts," *IEEE Trans. Inf. Forensics Secur.*, 7(3):1003-1017, 2012.
- [8] P. Ferrara, T. Bianchi, A.D. Rosa, A. Piva, "Image forgery localization via fine-grained analysis of CFA artifacts," *IEEE Trans. on Inf. Forensics Secur.*, 7(5):1566-1577, 2012.

- [9] I. Yerushalmy, H. Hel-Or, "Digital image forgery detection based on lens and sensor aberration," *Int. J. Comput. Vis.*, 92(1):71-91, 2011.
- [10] M. Chen, J. Fridrich, M. Goljan, J. Lukáš, "Determining image origin and integrity using sensor noise," *IEEE Trans. Inf. Forensics Secur.*, 3(1):74-90, 2008.
- [11] N. Rahmouni, V. Nozick, J. Yamagishi, I. Echizen, "Distinguishing computer graphics from natural images using convolution neural networks," in *Proc. IEEE Int. Workshop Inf. Forensics Secur.*, pp.1-6, 2017.
- [12] I.J. Yu, D.G. Kim, J.S. Park, J.U. Hou, S.H. Choi, H.K. Lee, "Identifying photorealistic computer graphics using convolutional neural networks," in *Proc. 24th IEEE Int. Conf. Image Process.*, pp.4093-4097, 2017.
- [13] W.Z. Quan, K. Wang, D.M. Yan, X.P. Zhang, "Distinguishing between natural and computer-generated images using convolutional neural networks," *IEEE Trans. Inf. Forensics Secur.*, 13(11):2772-2787, 2018.
- [14] S. McCloskey, M. Albright, "Detecting GAN-generated imagery using color cues," *arXiv preprint:1812.08247*, 2018.
- [15] F. Marra, D. Gagnaniello, L. Verdoliva, G. Poggi, "Do GANs leave artificial fingerprints," *arXiv preprint:1812.11842v1*, 2018.
- [16] H.D. Li, B. Li, S.Q. Tan, J.W. Huang, "Identification of deep network generated images using disparities in color components," *Signal Process.*, 174:107616, 2020.
- [17] D. Cozzolino, J. Thies, A. Rössler, C. Riess, "ForensicTransfer: weakly-supervised domain adaptation for forgery detection," *arXiv preprint:1812.02510*, 2018.
- [18] N. Yu, L. Davis, M. Fritz, "Attributing fake images to GANs: analyzing fingerprints in generated images," *arXiv preprint:1811.08180*, 2018.
- [19] D. Afchar, V. Nozick, J. Yamagishi, I. Echizen, "MesoNet: a compact facial video forgery detection network," *Proc. IEEE Int. Workshop Inf. Forensics Secur.*, pp.1-6, 2018.
- [20] H.H. Nguyen, F.M. Fang, J. Yamagishi, I. Echizen, "Multi-task Learning for Detecting and Segmenting Manipulated Facial Images and Videos," *Proc. 10th IEEE Int. Conf. Biom. Theory, Appl. Syst. (BTAS)*, pp.1-8, 2019.
- [21] K. Simonyan, A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. 32nd Int. Conf. Mach. Learn.*, pp.1-14, 2015.
- [22] D.P. Kingma, J.L. Ba, "ADAMM: a method for stochastic optimization," in *Proc. 3rd Int. Conf. Learn. Represent.*, pp.1-15, 2015.
- [23] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, 15(1):1929-1958, 2014.
- [24] T.T. Ng, S.F. Chang, J. Hsu, and M. Pepeljugoski, "Columbia photographic images and photorealistic computer graphics dataset," *ADVENT Tech. Rep. #205-2004-5*, Columbia University, 2005.
- [25] T. Karras, T. Aila, S. Laine, J. Lehtinen, "Progressive growing of GANs for improved quality, stability, and variation," *Proc. 6th Int. Conf. Learn. Represent.*, pp.1-26, 2018.
- [26] F. Yu, Y.D. Zhang, S.R. Song, A. Seff, J.X. Xiao, "LSUN: Construction of a large-scale image dataset using deep learning with humans in the loop," *arXiv preprint:1506.03365*, 2015.
- [27] Z.W. Liu, P. Luo, X.G. Wang, X.O. Tang, "Deep learning face attributes in the wild," *IEEE Int. Conf. on Comput. Vis.*, pp.3730-3738, 2015.