

DeeperForensics Challenge 2020 on Real-World Face Forgery Detection: Methods and Results

Liming Jiang, Zhengkui Guo, Wayne Wu, Zhaoyang Liu, Ziwei Liu, Chen Change Loy, Shuo Yang, Yuanjun Xiong, Wei Xia, Baoying Chen, Peiyu Zhuang, Sili Li, Shen Chen, Taiping Yao, Shouhong Ding, Jilin Li, Feiyue Huang, Liujuan Cao, Rongrong Ji, Changlei Lu, Ganchao Tan

Abstract

This paper reports methods and results in the DeeperForensics Challenge 2020 on real-world face forgery detection. The challenge employs the DeeperForensics-1.0 dataset, one of the most extensive publicly available real-world face forgery detection datasets, with 60,000 videos constituted by a total of 17.6 million frames. The model evaluation is conducted online on a high-quality hidden test set with multiple sources and diverse distortions. A total of 115 participants registered for the competition, and 25 teams made valid submissions. We will summarize the winning solutions and present some discussions on potential research directions.

1. Introduction

Recent years have witnessed exciting progress [1, 23, 4, 5, 20, 17] in automatic face swapping. Indeed, these techniques have eschewed the cumbersome hand-crafted face manipulation processes, hence facilitating the development of various popular softwares for face editing. From another perspective, these easy-to-access softwares, named “Deepfakes”, have also brought risks for being misused and

-
- Liming Jiang, Ziwei Liu, and Chen Change Loy are with S-Lab, Nanyang Technological University.
 - Zhengkui Guo is with The Chinese University of Hong Kong.
 - Wayne Wu and Zhaoyang Liu are with SenseTime Research.
 - Shuo Yang, Yuanjun Xiong, and Wei Xia are with Amazon Web Services.
 - Baoying Chen, Peiyu Zhuang, and Sili Li are with Shenzhen Key Laboratory of Media Information Content Security, Shenzhen University.
 - Shen Chen, Liujuan Cao, and Rongrong Ji are with Media Analytics and Computing Lab, Xiamen University.
 - Taiping Yao, Shouhong Ding, Jilin Li, and Feiyue Huang are with YouTu Lab, Tencent, Shanghai.
 - Changlei Lu and Ganchao Tan are with University of Science and Technology of China.

spread. Tampered videos on the internet could lead to possible perilous consequences, entailing the potential legitimate concerns among the general public and authorities. Therefore, effective face forgery detection methods become an urgent need to safeguard against these photorealistic fake videos, particularly in real-world scenarios where the video sources and distortions are unknown.

We organize the DeeperForensics Challenge 2020 with the aim to advance the state-of-the-art in face forgery detection. Participants are expected to develop robust and generic methods for forgery detection in real-world scenarios. The challenge uses DeeperForensics-1.0 [17], a large-scale real-world face forgery detection dataset that contains 60,000 videos with a total of 17.6 million frames¹. All source videos in DeeperForensics-1.0 are carefully collected, and fake videos are generated by a newly proposed end-to-end face swapping framework. Extensive real-world perturbations are applied to obtain a more challenging benchmark of larger scale and higher diversity. The dataset also features a hidden test set, which is richer in distribution than the publicly available training set, suggesting a better setting to simulate real-world scenarios. Besides, the hidden test set will be continuously updated to get future versions along with the development of Deepfakes technology. The evaluation of the challenge is performed online on the current version of the hidden test set.

In the following sections, we will describe the DeeperForensics Challenge 2020, summarize the winning solutions and results, and provide discussions to take a closer look at the current status and possible future development of real-world face forgery detection.

2. About the Challenge

2.1. Platform

The DeeperForensics Challenge 2020 is hosted on the CodaLab platform² in conjunction with ECCV 2020, The

¹ Project page: <https://liming-jiang.com/projects/DrF1/DrF1.html>.

² Challenge website: <https://competitions.codalab.org/competitions/25228>.

2nd Workshop on Sensing, Understanding and Synthesizing Humans³. The online evaluation is conducted using Amazon Web Services (AWS)⁴. First, participants register their teams on the CodaLab challenge website. Then, they are requested to submit their models to the AWS evaluation server (with one 16 GB Tesla V100 GPU for each team) to perform the online evaluation on the hidden test set. When the evaluation is done, participants receive the encrypted prediction files through an automatic email. Finally, they submit the result file to the CodaLab challenge website.

2.2. Dataset

The DeeperForensics Challenge 2020 employs the DeeperForensics-1.0 dataset [17] that was proposed in CVPR 2020. DeeperForensics-1.0 contains 60,000 videos constituted by a total of 17.6 million frames. The dataset features three appealing properties: good quality, large scale, and high diversity.

To ensure good quality, extensive data collection is conducted. The high-resolution (1920 × 1080) source videos are collected from 100 paid actors with four typical skin tones across 26 countries. Their eight expressions (*i.e.*, neutral, angry, happy, sad, surprise, contempt, disgust, fear) are recorded under nine lighting conditions by seven cameras at different locations. We further ask the actors to perform 53 supplementary expressions defined by 3DMM blendshapes [9] to make the dataset more diverse. Besides, a robust end-to-end face swapping framework, DF-VAE, is developed to generate the fake videos. In addition, seven types of real-world perturbations at five intensity levels are applied to obtain a more challenging benchmark of larger scale and higher diversity. Readers are referred to [17] for details.

An indispensable component of DeeperForensics-1.0 is the hidden test set, which is richer in distribution than the publicly available training set. The hidden test set suggests a better real-world face forgery detection setting: **1)** Multiple sources. Fake videos in-the-wild should be manipulated by different unknown methods; **2)** High quality. Threatening fake videos should have high quality to deceive human eyes; **3)** Diverse distortions. Different perturbations should be considered. The hidden test set will evolve by including more challenging samples along with the development of Deepfakes technology. The evaluation of the challenge is performed on its current version.

All the participants using the DeeperForensics-1.0 dataset should agree to its Terms of Use [7]. They are recommended but not restricted to train their algorithms on DeeperForensics-1.0. The use of any external datasets should be disclosed and follow the Terms of Use.

³ Workshop website: <https://sense-human.github.io>.

⁴ Online evaluation website: <https://aws.amazon.com>.

2.3. Evaluation Metric

Similar to Deepfake Detection Challenge (DFDC) [2], the DeeperForensics Challenge 2020 uses the binary cross-entropy loss (BCELoss) to evaluate the performance of face forgery detection models:

$$\text{BCELoss} = -\frac{1}{N} \sum_{i=1}^N [y_i \cdot \log(p(y_i)) + (1 - y_i) \cdot \log(1 - p(y_i))],$$

where N is the number of videos in the hidden test set, y_i denotes the ground truth label of video i (fake: 1, real: 0), and $p(y_i)$ indicates the predicted probability that video i is fake. A smaller BCELoss score is better, which directly contributes to a higher ranking. If the BCELoss score is the same, the one with less runtime will achieve a higher ranking. To avoid an infinite BCELoss that is both too confident and wrong, the score is bounded by a threshold value.

2.4. Timeline

The DeeperForensics Challenge 2020 lasted for nine weeks – eight weeks for the *development phase* and one week for the *final test phase*.

The challenge officially started at the ECCV 2020 Sense-Human Workshop on August 28, 2020, and it immediately entered the development phase. In the development phase, the evaluation is performed on the *test-dev* hidden test set, which contains 1,000 videos representing general circumstances of the full hidden test set. The *test-dev* hidden test set is used to maintain a public leaderboard. Participants can conduct four online evaluations (each with 2.5 hours of runtime limit) per week.

The final test phase started on October 24, 2020. The evaluation is conducted on the *test-final* hidden test set, containing 3,000 videos (also including test-dev videos) with a similar distribution as test-dev, for the final competition results. A total of two online evaluations (each with 7.5 hours of runtime limit) are allowed. The final test phase ended on October 31, 2020.

Finally, the challenge results were announced in December 2020. In total, 115 participants registered for the competition, and 25 teams made valid submissions.

3. Results and Solutions

Table 1. Final results of the top-5 teams in the DeeperForensics Challenge 2020. The runtime is shown in seconds.

Ranking	TeamName	UserName	BCELoss↓	Runtime↓
1	Forensics	BokingChen	0.2674	7690
2	RealFace	Iverson	0.3699	11368
3	VISG	zz110	0.4060	11012
4	jiashangplus	jiashangplus	0.4064	16389
5	Miao	miaotao	0.4132	19823

Among the 25 teams who made valid submissions, many participants achieve promising results. We show the final

results of the top-5 teams in Table 1. In the following sections, we will present the winning solutions of top-3 entries.

3.1. Solution of First Place

Team members: Baoying Chen, Peiyu Zhuang, Sili Li




Figure 1. The framework of the first-place solution.

As shown in Figure 1, the method designed by the champion team contains three stages, namely Face Extraction, Classification, and Output.

Face Extraction. They first extract 15 frames from each video at equal intervals using VideoCapture of OpenCV. Then, they use the face detector MTCNN [30] to detect the face region of each frame and expand the region by 1.2 times to crop the face image.

Classification. They define the prediction of the probability that the face is fake as the face score. They use EfficientNet [25] as the backbone, which was proven effective in the Deepfake Detection Challenge (DFDC) [2]. The results of three models (EfficientNet-B0, EfficientNet-B1 and EfficientNet-B2) are ensemble for each face.

Output. The final output score of a video is the predicted probability that the video is fake, which is calculated by the average of face scores for the extracted frames.

Implementation Details. The team employs EfficientNet pre-trained on ImageNet as the backbone. They select EfficientNet-B0, EfficientNet-B1, and EfficientNet-B2 for the model ensemble. In addition to DeeperForensics-1.0, they use some other public datasets, *i.e.*, UADFV [28], Deep Fake Detection [8], FaceForensics++ [24], Celeb-DF [29], and DFDC Preview [15]. They balance the class samples with the down-sampling mode. The code of the champion solution has been made publicly available⁵.

• **Training:** Inspired by the DFDC winning solution, appropriate data augmentation could contribute to better results. As for the data augmentation, the champion team uses the perturbation implementation in DeeperForensics-1.0 [6] during training. They only apply the image-level distortions: color saturation change (CS), color contrast change (CC), local block-wise (BW), white Gaussian noise in color components (GNC), Gaussian blur (GB) and JPEG compression (JPEG). They randomly mixup these distortions with a probability of 0.2. Besides, they also try other data augmentation [3], but the performance improvement is slim. The images are resized to 224×224 . The batch size is 128,

⁵ <https://github.com/beibuwandeluori/DeeperForensicsChallengeSolution>.

and the total training epoch is 50. They use AdamW optimizer [22] with initial learning rate of 0.001. Label smoothing is applied with a smoothing factor of 0.05.

• **Testing:** The testing pipeline follows the three stages in Figure 1. They clip the prediction score of each video in a range of $[0.01, 0.99]$ to reduce the large loss caused by the prediction errors. In addition to the best BCELoss score, their fastest execution speed may be attributed to the use of the faster face extractor MTCNN and the ensemble of three image-level models with fewer parameters.

3.2. Solution of Second Place

Team members: Shen Chen, Taiping Yao, Shouhong Ding, Jilin Li, Feiyue Huang, Liujuan Cao, Rongrong Ji




Figure 2. The framework of the second-place solution.

Face manipulated video contains two types of forgery traces, *i.e.*, image-level artifacts and video-level artifacts. The former refers to the artifacts such as blending boundaries and abnormal textures within image, while the latter is the face jitter problem between video frames. Most previous works only focused on artifacts in a specific modality and lacked consideration of both. The team in the second place proposes to use an attention mechanism to fuse the temporal information in videos, and further combine it with an image model to achieve better results.

The overall framework of their method is shown in Figure 2. First, they use RetinaFace [13] with 20% margin to detect faces in video frames. Then, the face sequence is fed into an image-based model and a video-based model, where the backbones are both EfficientNet-b5 [25] with NoisyStudent [27] pre-trained weights. The image-based model predicts frame by frame and takes the median of probabilities as the prediction. The video-based model takes the entire face sequence as the input and adopts an attention module to fuse the temporal information between frames. Finally, the per-video prediction score is obtained by averaging the probabilities predicted by the above two models.

Implementation Details. The team implements the proposed method via PyTorch. All the models are trained on 8 NVIDIA Tesla V100 GPUs. In addition to the DeeperForensics-1.0 dataset, they use three external datasets, *i.e.*, FaceForensics++ [24], Celeb-DF [29], and Diverse Fake Face Dataset [12]. They used the official splits provided by the above datasets to construct the training, validation and test sets. They balance the positive and negative samples through the down-sampling technique.

- **Training:** The second-place team uses the following data augmentations: RandAugment [11], patch Gaussian [21], Gaussian blur, image compression, random flip, random crop and random brightness contrast. They also employ the perturbation implementation in DeeperForensics-1.0 [6]. For the image-based model, they train a classifier based on EfficientNet-b5 [25], using binary cross-entropy loss as the loss function. They adopt a two-stage training strategy for the video-based model. In stage-1, they train an image-based classifier based on EfficientNet-b5. In stage-2, they fix the model parameters trained in stage-1 to serve as face feature extractor, and introduce an attention module to learn temporal information via nonlinear transformations and *softmax* operations. The input of the network is the face sequence (*i.e.*, 5 frames per video) in stage-2, and only the attention module and classification layers are trained. The binary cross-entropy loss is adopted as the loss function. The input size is scaled to 320×320 . Adam optimizer [19] is used with a learning rate of 0.0002, $\beta_1 = 0.9$, $\beta_2 = 0.999$, and weight decay of 0.00001. The batch size is 32. The total number of training epochs is set to 20, and the learning rate is halved every 5 epochs.

- **Testing:** They sample 10 frames at equal intervals for each video and detect faces by RetinaFace [13] as in the training phase. Then, the face images are resized to 320×320 . Test-time augmentation (TTA) (*e.g.*, flip) is applied to get 20 images (10 original and 10 flipped), which are fed into network to get the prediction score. They clip the prediction score of each video to $[0.01, 0.99]$ to avoid excessive losses on extreme error samples.

3.3. Solution of Third Place

Team members: Changlei Lu, Ganchoo Tan

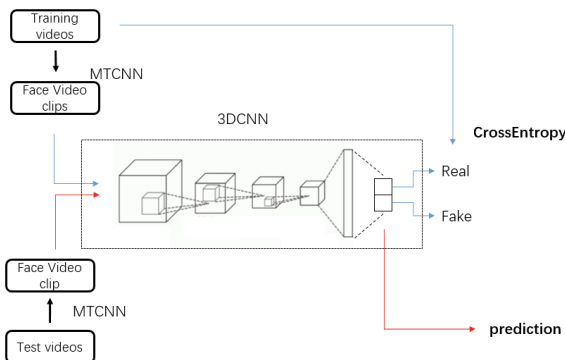


Figure 3. The framework of the third-place solution.

Similar to the second-place entry, the team in the third place also utilize the poor temporal consistency in existing face manipulation techniques. To this end, they propose to use a 3D convolutional neural network (3DCNN) to capture spatial-temporal features for forgery detection. The framework of their method is shown in Figure 3.

Implementation Details. First, the team crops faces in the video frames using the MTCNN [30] face detector. They combine all the cropped face images into a face video clip. Each video clip is then resized to $64 \times 224 \times 224$ or $64 \times 112 \times 112$. Various data augmentations are applied, including Gaussian blur, white Gaussian noise in color components, random crop, random flip, *etc.* Then, they use the processed video clips as the input to train a 3D convolutional neural network (3DCNN) using the cross-entropy loss. They examine three kinds of networks, I3D [10], 3D ResNet [16] and R(2+1)D [26]. These models are pre-trained on the action recognition datasets, *e.g.*, kinetics [18]. In addition to DeeperForensics-1.0, they use three external public face manipulation datasets, *i.e.*, the DFDC dataset [14], Deep Fake Detection [8], and FaceForensics++ [24].

4. Discussion

The methods mentioned above have considered different potential aspects in developing a robust face forgery detection model. We are glad to find the winning solutions achieve promising performance in the DeeperForensics Challenge 2020. In summary, there are three key points inspired by these methods that could improve real-world face forgery detection. 1) Strong backbone. Backbone selection of the forgery detection models is important. The high-performance winning solutions are based on state-of-the-art EfficientNet. 2) Diverse augmentations. Applying appropriate data augmentations may better simulate real-world scenarios and boost the model performance. 3) Temporal information. Since the primary detection target is the fake videos, temporal information can be a critical clue to distinguish the real from the fake.

Despite the promising results, we believe that there is still much room for improvement in the real-world face forgery detection task. 1) More suitable and diverse data augmentations may contribute to a better simulation of real-world data distribution. 2) Developing a robust detection method that can cope with unseen manipulation methods and distortions is a critical problem. At this stage, we observe that the model training is data-dependent. Although data augmentations can help improve the performance to a certain extent, the generalization ability of most forgery detection models is still poor. 3) Different artifacts in the Deepfakes videos (*e.g.*, checkerboard artifacts, fusion boundary artifacts) remain rarely explored.

Acknowledgments. We thank Amazon Web Services for sponsoring the prize of this challenge. The organization of this challenge is also supported by A*STAR through the Industry Alignment Fund - Industry Collaboration Projects Grant.

References

- [1] DeepFaceLab. <https://github.com/iperov/DeepFaceLab>. Accessed: 2019-08-20. 1
- [2] Deepfake Detection Challenge. <https://www.kaggle.com/c/deepfake-detection-challenge>. Accessed: 2020-02-15. 2, 3
- [3] Deepfake detection (DFDC) solution by selimsef. https://github.com/selimsef/dfdc_deepfake_challenge. Accessed: 2020-10-30. 3
- [4] DeepFakes. <https://github.com/deepfakes/faceswap>. Accessed: 2019-08-16. 1
- [5] faceswap-GAN. <https://github.com/shaoanlu/faceswap-GAN>. Accessed: 2019-08-16. 1
- [6] Perturbation implementation in DeeperForensics-1.0. <https://github.com/EndlessSora/DeeperForensics-1.0/tree/master/perturbation>. Accessed: 2020-10-30. 3, 4
- [7] Terms of use: DeeperForensics-1.0 dataset. https://github.com/EndlessSora/DeeperForensics-1.0/blob/master/dataset/Terms_of_Use.pdf. Accessed: 2020-05-21. 2
- [8] Google AI Blog. Contributing data to deepfake detection research. <https://ai.googleblog.com/2019/09/contributing-data-to-deepfake-detection.html>. Accessed: 2019-09-25. 3, 4
- [9] Chen Cao, Yanlin Weng, Shun Zhou, Yiyong Tong, and Kun Zhou. FaceWarehouse: A 3d facial expression database for visual computing. *TVCG*, 20:413–425, 2013. 2
- [10] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017. 4
- [11] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. RandAugment: Practical automated data augmentation with a reduced search space. In *CVPRW*, 2020. 4
- [12] Hao Dang, Feng Liu, Joel Stehouwer, Xiaoming Liu, and Anil K Jain. On the detection of digital face manipulation. In *CVPR*, 2020. 3
- [13] Jiansheng Deng, Jia Guo, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. RetinaFace: Single-shot multi-level face localisation in the wild. In *CVPR*, 2020. 3, 4
- [14] Brian Dolhansky, Joanna Bitton, Ben Pflaum, Jikuo Lu, Russ Howes, Menglin Wang, and Cristian Canton Ferrer. The deepfake detection challenge dataset. *arXiv preprint*, arXiv:2006.07397, 2020. 4
- [15] Brian Dolhansky, Russ Howes, Ben Pflaum, Nicole Baram, and Cristian Canton Ferrer. The deepfake detection challenge (DFDC) preview dataset. *arXiv preprint*, arXiv:1910.08854, 2019. 3
- [16] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Learning spatio-temporal features with 3D residual networks for action recognition. In *ICCVW*, 2017. 4
- [17] Liming Jiang, Ren Li, Wayne Wu, Chen Qian, and Chen Change Loy. DeeperForensics-1.0: A large-scale dataset for real-world face forgery detection. In *CVPR*, 2020. 1, 2
- [18] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint*, arXiv:1705.06950, 2017. 4
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*, arXiv:1412.6980, 2014. 4
- [20] Lingzhi Li, Jianmin Bao, Hao Yang, Dong Chen, and Fang Wen. FaceShifter: Towards high fidelity and occlusion aware face swapping. In *CVPR*, 2020. 1
- [21] Raphael Gontijo Lopes, Dong Yin, Ben Poole, Justin Gilmer, and Ekin D Cubuk. Improving robustness without sacrificing accuracy with patch gaussian augmentation. *arXiv preprint*, arXiv:1906.02611, 2019. 4
- [22] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. 3
- [23] Ivan Petrov, Daiheng Gao, Nikolay Chervoniy, Kunlin Liu, Sugasa Marangonda, Chris Umé, Jian Jiang, Luis RP, Sheng Zhang, Pingyu Wu, et al. DeepFaceLab: A simple, flexible and extensible face swapping framework. *arXiv preprint*, arXiv:2005.05535, 2020. 1
- [24] Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. FaceForensics++: Learning to detect manipulated facial images. In *ICCV*, 2019. 3, 4
- [25] Mingxing Tan and Quoc Le. EfficientNet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019. 3, 4
- [26] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, 2018. 4
- [27] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves ImageNet classification. In *CVPR*, 2020. 3
- [28] Xin Yang, Yuezun Li, and Siwei Lyu. Exposing deep fakes using inconsistent head poses. In *ICASSP*, 2019. 3
- [29] Pu Sun Honggang Qi Yuezun Li, Xin Yang and Siwei Lyu. Celeb-DF: A large-scale challenging dataset for deepfake forensics. In *CVPR*, 2020. 3
- [30] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *SPL*, 23:1499–1503, 2016. 3, 4