

Recurrent Convolutional Strategies for Face Manipulation Detection in Videos

Ekraam Sabir, Jiaxin Cheng, Ayush Jaiswal, Wael AbdAlmageed, Iacopo Masi, Prem Natarajan
USC Information Sciences Institute, Marina del Rey, CA, USA

{esabir, chengjia, ajaiswal, wamageed, iacopo, pnataraj}@isi.edu

Abstract

The spread of misinformation through synthetically generated yet realistic images and videos has become a significant problem, calling for robust manipulation detection methods. Despite the predominant effort of detecting face manipulation in still images, less attention has been paid to the identification of tampered faces in videos by taking advantage of the temporal information present in the stream. Recurrent convolutional models are a class of deep learning models which have proven effective at exploiting the temporal information from image streams across domains. We thereby distill the best strategy for combining variations in these models along with domain specific face preprocessing techniques through extensive experimentation to obtain state-of-the-art performance on publicly available video-based facial manipulation benchmarks. Specifically, we attempt to detect Deepfake, Face2Face and FaceSwap tampered faces in video streams. Evaluation is performed on the recently introduced FaceForensics++ dataset, improving the previous state-of-the-art by up to 4.55% in accuracy.

1. Introduction

A spate of recent incidents have increased the scrutiny of online misinformation [11, 5]. This has spurred research on both analysis and detection of misinformation [41, 35]. Misinformation can be manifested in different ways — direct manipulation of information or presentation of unmanipulated content in a misleading context. Digital image manipulation such as copy-move and splicing [45, 44] are examples of deliberate manipulation, while image repurposing [17, 36, 18] is an example of misleading context. Out of the wide range of manipulations on different modalities, *automatic* manipulation of digital content has recently gained attention. In particular, facial manipulations lately became very popular as a way for disseminating false information or even to libel celebrities or very well-known

people [2]. It is not surprising that faces are preferred over other objects, since the face is used everyday in our society as a means for human communication and associating information with identity.

Prior to the rise of deep learning, a malicious attacker would *manually* prepare counterfeit media either using Adobe Photoshop¹ or GIMP², making the process tedious, but nowadays the scenario has changed drastically: before machine learning came in the process, tools only aided the user in content creation, whilst currently, machine-learning-aided tools create content without manual intervention. The user is just required to “babysit” the training process, most of the time using easy-to-use Graphical User Interfaces (GUI) [3, 1].

More precisely, fueled by the recent success of Generative Adversarial Networks (GANs) [12] along with the availability of graphical processors (GPUs), any amateur user is capable of producing completely synthetic yet hyper-realistic content. The realism achieved by the aforementioned machine learning components is so high that even humans have difficulties in assessing whenever a face picture is naturally captured with a camera or, artificially produced. This statement is even more true for *single still images*. For instance, websites such as thispersondoesnotexist.com offer evidences of the level of realism reached by the aforementioned tools. The level of artefacts produced by such systems is so subtle that the only clues to assess if a face is real or fake are (i) subtle inconsistency in the hairs — too straight, with disconnected strands or simply unnatural, (ii) unnaturally asymmetric face, (iii) weird teeth, and more importantly and most of the time, (iv) other more clear inconsistencies are not localized on the face yet in the background. These latter artefacts are peculiar of generative networks that render the entire head along with the background e.g., StarGAN [7].

While we believe that it is important to develop detectors to pick up those salient irregularities in still images [49, 4] to stem the spread of false information, additional, orthogonal

¹www.adobe.com/products/photoshopfamily.html

²www.gimp.org

important features for detecting manipulations in videos are captured by the temporal coherence inherently present in a video stream. In this sense, face manipulation in videos has recently gained interest upon still images. The increased proliferation of fake videos may be attributed to two reasons:

1. Transposing a person’s identity or expression with somebody else’s, is easier now given the off-the-self availability of machine-learning-aided face swapping or face reenactment tools [39, 1, 31, 3].
2. A video is more likely to be believable rather than a still image since it demonstrates the activity in progress and also, in principle, it requires a much painstaking effort in manipulating *coherently* all the frames.

In light of above observations and considering that face manipulation generation tools [39, 1, 31, 3] do not enforce temporal coherence in the synthesis process and perform manipulations on a frame-by-frame basis, we propose to leverage temporal artefacts as a means for indication of abnormal faces in a video stream. Unlike state-of-the-art fake detection methods [49, 4], we explore recurrent convolutional models to exploit temporal discrepancies for improving upon the current practice. To remove other confounding factors related to the rigid motion of the face in a video, we also explore face alignment methods. We optimize a deep learning model architecture over these two factors which gives state-of-the-art performance in face manipulation detection accuracy.

The rest of the manuscript is organized as follows: Section 2 presents the related work on temporal processing with deep models, recent face manipulation benchmarks along with manipulation detection techniques recently published. Section 3 explains the proposed methodology comprising of face preprocessing, feature extraction and finally prediction based on a bi-direction recurrent model. The experimental evaluation is presented in Section 4 with quantitative results on FaceForensics++ (FF++) [34]. Section 5 concludes the paper.

2. Related Work

Video processing with deep models. Activity recognition in videos has well developed literature and can be used to gain insights into processing videos, taking advance of temporal information. There are three major approaches in this area. The first line of research develops from a two stream network [37], where an RGB video frame and its optical flow version are processed in two separate branches in the network followed by a fusion mechanism. This former strategy of video processing aims to capture temporal

information and motion across frames with the usage of optic flow. The second is a single stream network supported by recurrent convolutional layers: perceptual knowledge of the content in each frame is gathered using a separate convolutional neural network (CNN) that extracts high-level semantic features, while a recurrent model is trained on top of those features to perform decisions over the temporal dimension [10]. The third line of development resides in 3D convolutions [40, 19] as a local building block in the network to learn rich spatio-temporal features. With respect to all the previous mentioned strategies, we argue that methods based on two-stream architecture are effective in action recognition but not relevant to capture the subtle flickering artefacts that a generator may produce in a video; on the other hand, 3D convolutions could be better suited for this purpose but they greatly increase the number of learned filters. For all these reasons, we use a main backbone CNN encoder and capture temporal anomaly in face appearance using a recurrent model. Recurrent models have been widely used in computer vision for face landmark detection [26], face age progression [42] and even face parsing [27]; nevertheless, to the best of our knowledge they have not been employed before for video manipulation, with the only exception of [13].

Face manipulation benchmarks. Regrettably, compelling datasets for face manipulation detection and evaluation in videos had been lacking in the community; some previous attempts [49] generated face swapped images using an iOS app and an open-source face swap software using still images; though Zhou *et al.* used this set for evaluation, it did not provide manipulation for video streams.

Video-based face manipulation became available for the community with the recent release of FaceForensics [33], followed by its extended and improved version called FaceForensics++ [34]. FaceForensics (FF) released Face2Face manipulated videos [39]. FaceForensics++ (FF++) is an extension of FF, further augmenting the collection with *Deepfake* [1] and *FaceSwap* [28] manipulations. The set comprises of 1,000 videos organized into a single split where 720 videos are reserved for training and 140 videos used for validation and the same amount for test. All videos are collected from Youtube and capture “talking heads” or anchor men/women presenting news. In general, the subject faces the camera and cooperates with it.

Regarding the manipulation types, *Deepfakes* are generated via a two-branch autoencoder network. Deepfake system is an identity-swapping method and trained given two subjects. In this base version, a single autoencoder needs to be trained for a pair of individuals. The autoencoder shares a single encoder for compressing the information while two decoders reconstruct each one of the subject’s images. At test time, the two decoders are inverted to obtain the final face identity manipulation also known as face swapping; on

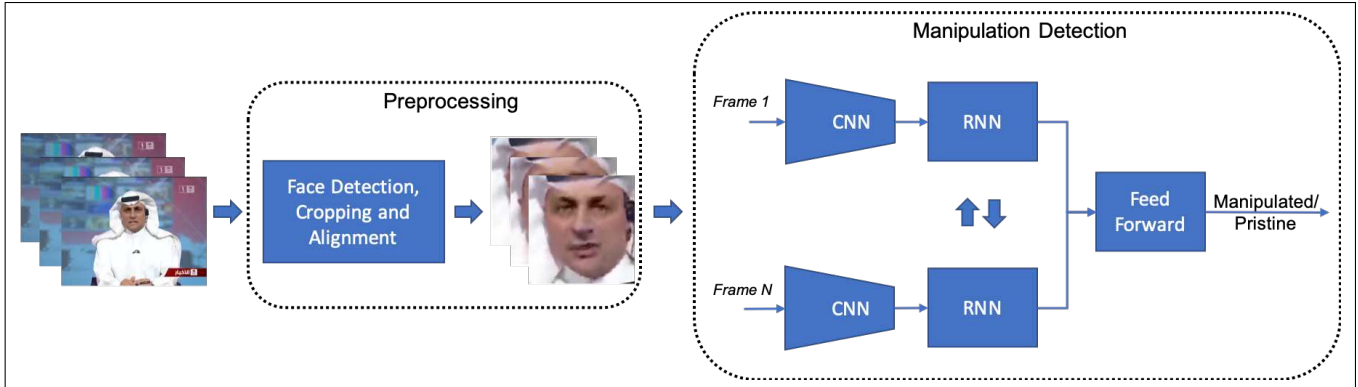


Figure 1: The overall pipeline is a two step process. The first step detects, crops and aligns faces on a sequence of frames. The second step is manipulation detection with our recurrent convolutional model.

the other hand, *FaceSwap* is a graphics-based approach to attain the same objective of swapping the identity of subjects. Unlike *Deepfakes*, *FaceSwap* can be applied to an unlimited number of subjects though is more prone to severe artefacts if some of its inside modules fail. Finally, *Face2Face* [39] is a graphics-based approach yet achieves very realistic facial reenactment. Reenactment transfers the expression and pose of a source character into a target video, while the identity of the target subject remains unvaried: i.e. it offers the same functionality of animating a pre-recorded video. For more details, we refer to [34].

Face Manipulation Detection. Copy-move and splicing detection datasets and methods are abundant for images [32, 46, 45, 47] and less dominant in videos [20]. However, due to the recent emergence of face manipulation problem the literature is relatively sparse. Rossler *et al.* [34] introduced baselines implemented through existing methods and thus trained an XceptionNet [8] architecture to establish state-of-the-art on the FF+ benchmark. In the meantime, MesoNet [4] introduces two CNN based architectures for face manipulation detection: they take a mesoscopic approach to manipulation detection, which combines information from both low level (microscopic) and high level (macroscopic) features. Their main novelty resides in the MesoInception block that extends the Inception module [38] with the usage of dilated convolution [48]. Zhou *et al.* trained a GoogLeNet based architecture for detecting tampered face in still images, in synergy with another model working with steganalysis features and trained with triplet loss. The two models’ scores are combined through late fusion. In order to improve transfer learning of features Cozzolino *et al.* [9] leverage the abundant data present for face manipulation by using FF++ collections and adapted the model to newer domain using *ForensicTransfer*: an encoder-decoder architecture provided with an activation loss measures the activation of the latent space vector distinctively for the real class and the fake class. The closest

prior work to ours is [13]: unlike them, we experiment with face alignment as a means for removing confounding factors in detecting facial manipulations and we make use of bidirectional recurrency rather than just mono-directional. Additionally, we attempt to detect multiple types of face manipulation while the previous work is centered around deepfake detection.

3. Method

The overall approach for manipulation detection is a two step process: cropping and alignment of faces from video frames, followed by manipulation detection over the pre-processed facial region. We explain both the steps in detail in this section.

3.1. Face preprocessing

For cropping the face region, we use the masks provided by [34], generated using computer graphics [39]. Since it has been shown in [29, 30] that face alignment is beneficial for face recognition, we employ face alignment here as well. In particular, we have experimented with two techniques for alignment: (i) explicit alignment using facial landmarks, where a reference coordinate system and the tightness of the face crop are decided *a priori* and all the faces are aligned to this reference coordinate system so that any *rigid* motion of the face is compensated, and (ii) implicit alignment that uses a Spatial Transformer Network (STN) [16, 43] based on an affine transformation. In the latter case, the network predicts the alignment parameters, conditioned on the input image, thus may learn to “zoom” on particular parts of the face, as necessary to minimize the expected loss in the training set.

In both cases the core idea is that we want the recurrent convolutional model to take as input a face “tubelet” [22, 23], which is a sequence of spatio-temporal, tightly aligned face crops across video frames.

Manipulation	Frames	FF++ [34]	ResNet50	DenseNet	ResNet50 + Alignment	DenseNet + Alignment	ResNet50 + Alignment + BiDir	DenseNet + Alignment + BiDir
Deepfake	1	93.46	94.8	94.5	96.1	96.4	-	-
	5	-	94.6	94.7	96.0	96.7	94.9	96.9
Face2Face	1	89.8	90.25	90.65	89.31	87.18	-	-
	5	-	90.25	89.8	92.4	93.21	93.05	94.35
FaceSwap	1	92.72	91.34	91.04	93.85	96.1	-	-
	5	-	90.95	93.11	95.07	95.8	95.4	96.3

Table 1: Accuracy for manipulation detection across all manipulation types. DenseNet with alignment and bidirectional recurrent network is found to perform best. FF++ [34] is the baseline in these experiments.

Manipulation	Base	Variation	
		Spatial Transformer	Multi Recurrence
Deepfake	96.9	91.7	94.4
Face2Face	94.35	87.46	89.9
FaceSwap	96.3	93.2	94.8

Table 2: Results on using variations to the recurrent convolutional architecture. Both spatial-transformer networks and multi-recurrent networks exhibit a decline in performance.

Landmark-based alignment. Face images are aligned using a simple similarity transformation (four degrees of freedom), compensating for isotropic scale, in-plane rotation, and 2D translation. Most of the faces in the dataset are near-frontal and thus, it was sufficient to employ an accurate yet fast landmark detector method [24] implemented through [25]. Though [24] returns dense landmarks on the face, we select only a set of seven sparse points located on the most discriminative features of the faces (corners of the eyes, the tip of the nose, and corners of the mouth). Following the similarity transformation, faces are aligned with a loose crop at a 224×224 resolution.

Spatial Transformer Network. A spatial transformer network (STN) [16] performs spatial alignment of data with learnable affine transformation parameters. It can be inserted between feature maps in deep learning networks. It comprises three components: a localization net, a grid generator and a sampler. The localization net predicts the affine transformation parameters and the grid generator and sampler warp the input feature map with the affine parameters to produce the output feature maps. They have been shown to focus on spatially important areas of the input feature map for improved performance [16].

3.2. Video-based Face Manipulation Detection

For manipulation detection, we use a recurrent-convolutional network similar to [10, 13], where the input

is a sequence of frames from the query video. The intuition behind this model is to exploit temporal discrepancies across frames. Temporal discrepancies are expected to occur in images, since manipulations are performed on a frame-by-frame basis. As such, low level artifacts caused by manipulations on faces are expected to further manifest themselves as temporal artifacts with inconsistent features across frames. There are two differences from [10] in our implementation: (1) instead of using CaffeNet [21], we explore more suitable CNN architectures for the problem, and (2) instead of averaging recurrent features across all time-steps, we extract the final output of the recurrent network. Our work is also different from [13], in that we train our model end-to-end, whereas they use pre-trained CNNs. Fig. 1 shows the model diagram.

Backbone encoding network. In our experiments, we explore ResNet [14] and DenseNet [15] for the CNN component of the model. There are two reasons for exploring these CNNs. FaceForensics++ [34] is a low resource dataset with a 1,000 videos and to avoid overfitting, authors had to use pre-trained XceptionNet [8] with fixed feature extraction layers. For end-to-end trainability, we chose ResNet [14] which was shown to be easily trainable by the authors. Additionally, manipulation artifacts exhibit low level features (such as discontinuous jawlines, blurred eyes etc.) which do not require high level face semantic features. DenseNet is also a suitable CNN architecture, because it extracts features at different levels of hierarchy [15].

Regardless of the architecture used, the backbone network is firstly trained on the FF++ training split minimizing cross-entropy loss for binary classification to develop features to discern real faces from synthetics. The backbone is then extended with RNN and finally trained end-to-end under multiple strategies.

RNN training strategies. In our approach we experiment with the recurrent models placed at different locations of the backbone network: it connects together the backbone network to act as a feature learner that passes features to the RNN aggregating inputs over time. The final system

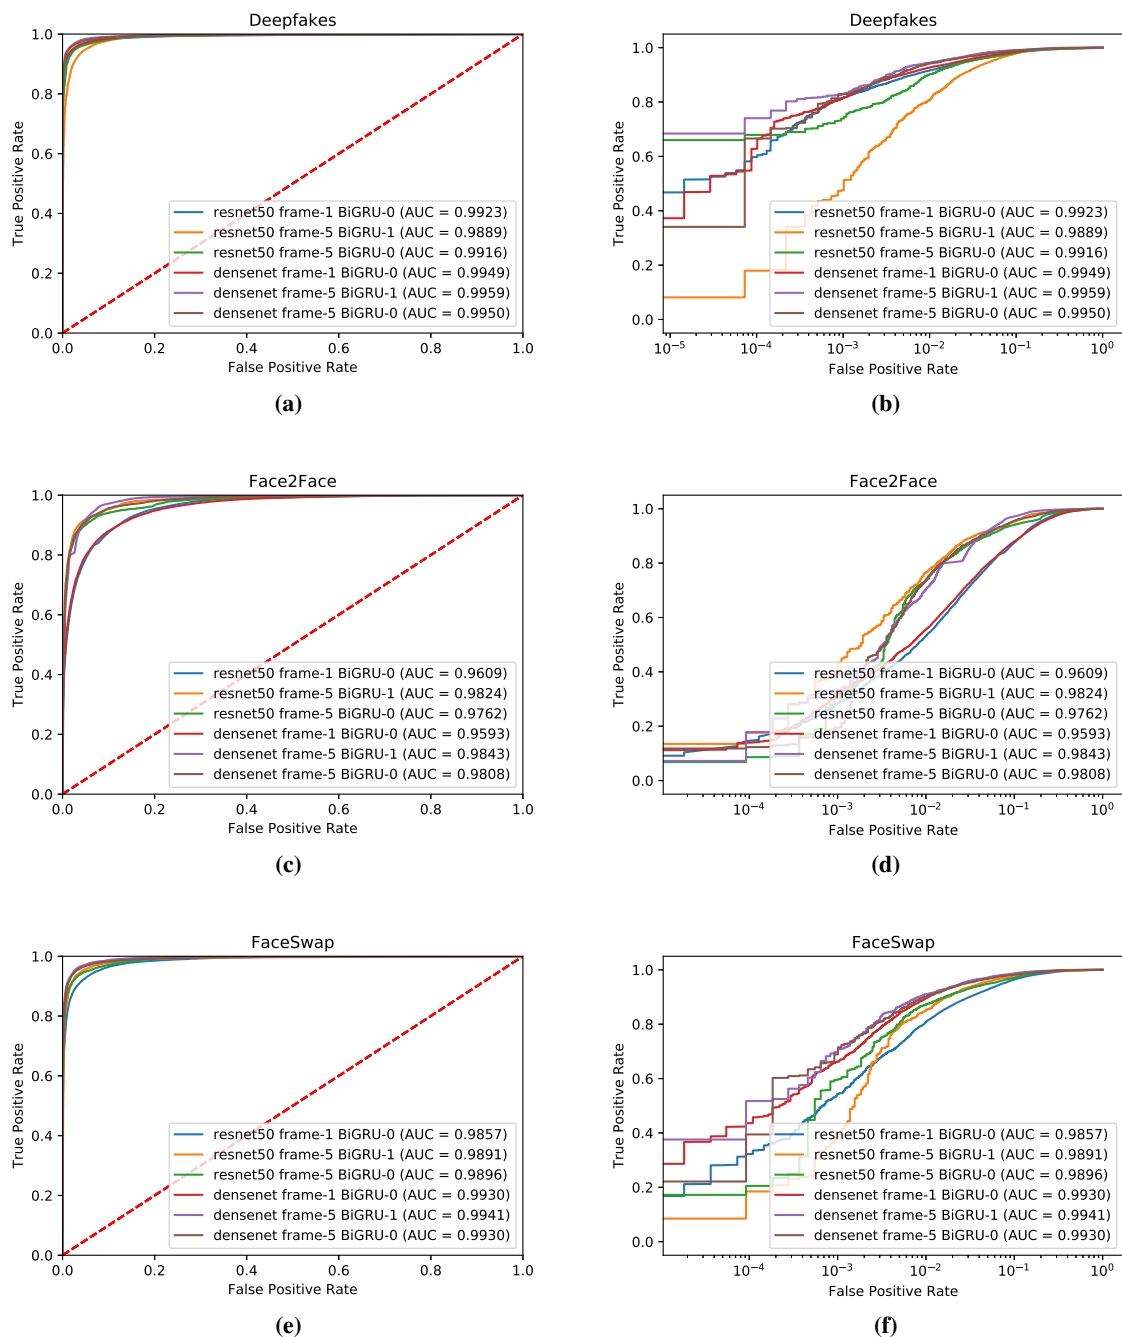


Figure 2: ROC plots for all manipulation types. Each row corresponds to a different manipulation type. The left column is a linear plot, while the right column has linear-log plots to better analyze the false alarm region. The legend reports ablation studies performed in our experiments: in particular, each experiment describes the backbone network used, how many frames are employed, if bidirectional GRU cells are used or not; finally each item offers the AUC.

is ultimately trained end-to-end. We experiment with two strategies: the first simply uses a single recurrent network on top of the final features from the backbone network; alternatively we attempt to learn multiple recurrent networks

at different level of the hierarchy of the backbone net: in [4], the authors propose a model that exploits features at a mesoscopic level. Their intuition is that purely microscopic and macroscopic features are not well suited for the

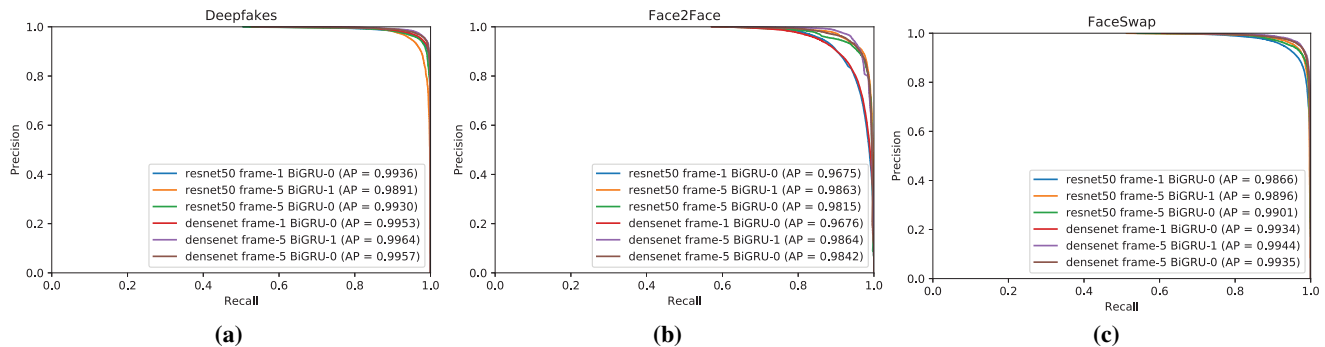


Figure 3: Precision Recall curves for all manipulation types. Each column corresponds to a different manipulation type. The legend is similar to Fig. 2. Each item also offers the average precision (AP) score.

face manipulation detection task. We incorporate this idea into our framework by extracting features at multiple feature levels from CNNs for manipulation detection. These features are processed in individual recurrent networks. We expect this new multi-recurrent-convolutional model to utilize micro, meso and macroscopic features for manipulation detection.

4. Experiments

Our evaluation metric is accuracy for a fair comparison with baselines in [34]. Additionally, we report area under the receiver operating curve (AUC) scores. All numbers are reported on FaceForensics++. For training, we use Adam optimizer with $1e-4$ learning rate. We use GRU cells [6] for our recurrent network. Additionally, all results are on the heavily compressed version of dataset from [34]. We do not evaluate on high and low quality videos since the baseline performance for those is already above 98% in [34].

Table 1 shows our results on *Deepfakes*, *Face2Face* and *FaceSwap* manipulation. These results show the impact of variations in architecture. We perform experiments to validate if face alignment improves performance. We use the landmark based face alignment discussed in Section 3.1. We also check if the temporal aspect of videos helps improve performance. We validate this with two variations: comparing five frame vs single frame input and uni- vs bi-directional recurrent network for five frame input experiments. We use a single recurrent network here on top of final CNN features. We consistently find (1) DenseNet to outperform ResNet, (2) face alignment to give improvement and (3) a sequence of images to be better than single frame input. We also find bidirectional recurrence to be superior to uni-directional recurrence. Fig. 2 and Fig. 3 report ROC and precision-recall (PR) curves for the same results. Since the false alarm region is hard to discern, we also show the linear-log plots for the ROC curve.

Table 2 shows our results from the choice of face align-

ment method and multiple levels of recurrence of our model. We use the best model from Table 1, namely the five-frame-bidirectional-densenet as our base model. The first variation is to use a spatial transformer network (STN) for learning an affine alignment of faces instead of a landmark based alignments as discussed in Section 3.1. For our experiments, we use a simple CNN with 2 convolution and max-pooling layers each followed by a feedforward network as the localization net and bilinear interpolation for the sampler. The second variation is the multi-recurrent model which involves extracting recurrent features at multiple levels of CNNs as explained in Section 3.2. Specifically, since the DenseNet used in our experiments has four blocks for generating feature maps, we build four recurrent networks for these. We find both strategies to be unsuccessful for improving performance. With respect to STN, we notice that training is relatively unstable and this may be due to changes in affine parameters during training which impacts the performance of the following CNN network. This may be a possible reason for performance drop when using STN. A possible reason for the relatively poor performance of the multi-recurrence model is the significant increase in the number of parameters, which leads to overfitting, given that the FaceForensics++ dataset has a limited number of samples (1,000 videos).

5. Conclusion

Misinformation in online content is increasing and there is an exigent need for detecting such content. Face manipulation in videos is one aspect of the larger problem. In this work we showed that a combination of recurrent-convolutional model and face alignment approach improves upon the state-of-the-art. We also explored different strategies for both alignment and combining CNN features through recurrence. We found a landmark based face alignment with bidirectional-recurrent-densenet to perform the best for face manipulation detection in videos.

Acknowledgments

This work is based on research sponsored by the Defense Advanced Research Projects Agency under agreement number FA8750-16-2-0204. 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 the Defense Advanced Research Projects Agency or the U.S. Government.

References

- [1] Deepfake project - non official project based on original deepfakes thread. Available at github.com/deepfakes/faceswap, January 2018. 1, 2
- [2] Deepfakes porn has serious consequences. www.bbc.com/news/technology-42912529, February 2018. 1
- [3] Deepfacelab project. Available at github.com/iperov/DeepFaceLab, April 2019. 1, 2
- [4] D. Afchar, V. Nozick, J. Yamagishi, and I. Echizen. MesoNet: a Compact Facial Video Forgery Detection Network. In *2018 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–7, Dec. 2018. 1, 2, 3, 5
- [5] H. Allcott and M. Gentzkow. Social Media and Fake News in the 2016 Election. *Journal of Economic Perspectives*, 31(2):211–236, May 2017. 1
- [6] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. *arXiv:1406.1078 [cs, stat]*, June 2014. arXiv: 1406.1078. 6
- [7] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *CVPR*, pages 8789–8797, 2018. 1
- [8] F. Chollet. Xception: Deep Learning With Depthwise Separable Convolutions. In *CVPR*, pages 1251–1258, 2017. 3, 4
- [9] D. Cozzolino, J. Thies, A. Rössler, C. Riess, M. Nießner, and L. Verdoliva. Forensicttransfer: Weakly-supervised domain adaptation for forgery detection. *arXiv preprint arXiv:1812.02510*, 2018. 3
- [10] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-Term Recurrent Convolutional Networks for Visual Recognition and Description. In *CVPR*, pages 2625–2634, 2015. 2, 4
- [11] E. Ferrara. Disinformation and Social Bot Operations in the Run Up to the 2017 French Presidential Election. SSRN Scholarly Paper ID 2995809, Social Science Research Network, Rochester, NY, June 2017. 1
- [12] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, pages 2672–2680, 2014. 1
- [13] D. Güera and E. J. Delp. Deepfake video detection using recurrent neural networks. In *AVSS*, pages 1–6. IEEE, 2018. 2, 3, 4
- [14] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *CVPR*, pages 770–778, 2016. 4
- [15] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger. Densely Connected Convolutional Networks. In *CVPR*, pages 4700–4708, 2017. 4
- [16] M. Jaderberg, K. Simonyan, A. Zisserman, et al. Spatial transformer networks. In *NIPS*, pages 2017–2025, 2015. 3, 4
- [17] A. Jaiswal, E. Sabir, W. AbdAlmageed, and P. Natarajan. Multimedia Semantic Integrity Assessment Using Joint Embedding Of Images And Text. In *ACMMM*, pages 1465–1471, New York, NY, USA, 2017. 1
- [18] A. Jaiswal, Y. Wu, W. AbdAlmageed, I. Masi, and P. Natarajan. AIRD: Adversarial Learning Framework for Image Repurposing Detection. In *CVPR*, 2019. 1
- [19] S. Ji, W. Xu, M. Yang, and K. Yu. 3d convolutional neural networks for human action recognition. *TPAMI*, 35(1):221–231, 2013. 2
- [20] S. Jia, Z. Xu, H. Wang, C. Feng, and T. Wang. Coarse-to-fine copy-move forgery detection for video forensics. *IEEE Access*, 6:25323–25335, 2018. 3
- [21] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 675–678. ACM, 2014. 4
- [22] V. Kalogeiton, P. Weinzaepfel, V. Ferrari, and C. Schmid. Action tubelet detector for spatio-temporal action localization. In *ICCV*, pages 4405–4413, 2017. 3
- [23] K. Kang, H. Li, T. Xiao, W. Ouyang, J. Yan, X. Liu, and X. Wang. Object detection in videos with tubelet proposal networks. In *CVPR*, pages 727–735, 2017. 3
- [24] V. Kazemi and J. Sullivan. One Millisecond Face Alignment with an Ensemble of Regression Trees. In *CVPR*, pages 1867–1874, 2014. 4
- [25] D. E. King. Dlib-ml: A Machine Learning Toolkit. *Journal of Machine Learning Research*, 10(Jul):1755–1758, 2009. 4
- [26] H. Lai, S. Xiao, Y. Pan, Z. Cui, J. Feng, C. Xu, J. Yin, and S. Yan. Deep recurrent regression for facial landmark detection. *IEEE Transactions on Circuits and Systems for Video Technology*, 28(5):1144–1157, 2018. 2
- [27] S. Liu, J. Shi, J. Liang, and M.-H. Yang. Face parsing via recurrent propagation. In *BMVC*, 2017. 2
- [28] Marek. 3d face swapping implemented in Python. Contribute to MarekKowalski/FaceSwap development by creating an account on GitHub, Apr. 2019. original-date: 2016-06-19T00:09:07Z. 2
- [29] I. Masi, F. Chang, J. Choi, S. Harel, J. Kim, K. Kim, J. Leksut, S. Rawls, Y. Wu, T. Hassner, W. AbdAlmageed, G. Medioni, L. Morency, P. Natarajan, and R. Nevatia. Learning pose-aware models for pose-invariant face recognition in the wild. *TPAMI*, 41(2):379–393, Feb 2019. 3

- [30] I. Masi, A. Tran, T. Hassner, G. Sahin, and G. Medioni. Face-specific data augmentation for unconstrained face recognition. *IJCV*, pages 1–26, 2019. 3
- [31] Y. Nirkin, I. Masi, A. Tran, T. Hassner, and G. Medioni. On face segmentation, face swapping, and face perception. In *AFGR*, pages 98–105, 2018. 2
- [32] R. C. Pandey, S. K. Singh, and K. K. Shukla. Passive copy-move forgery detection in videos. In *2014 International Conference on Computer and Communication Technology (IC-CCT)*, pages 301–306, Sept. 2014. 3
- [33] A. Rossler, D. Cozzolino, L. Verdoliva, C. Riess, J. Thies, and M. Niener. FaceForensics: A Large-scale Video Dataset for Forgery Detection in Human Faces. *arXiv:1803.09179 [cs]*, Mar. 2018. arXiv: 1803.09179. 2
- [34] A. Rossler, D. Cozzolino, L. Verdoliva, C. Riess, J. Thies, and M. Niener. FaceForensics++: Learning to Detect Manipulated Facial Images. *arXiv:1901.08971 [cs]*, Jan. 2019. arXiv: 1901.08971. 2, 3, 4, 6
- [35] N. Ruchansky, S. Seo, and Y. Liu. CSI: A Hybrid Deep Model for Fake News Detection. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM '17*, pages 797–806, New York, NY, USA, 2017. ACM. event-place: Singapore, Singapore. 1
- [36] E. Sabir, W. AbdAlmageed, Y. Wu, and P. Natarajan. Deep Multimodal Image-Repurposing Detection. In *Proceedings of the 26th ACM International Conference on Multimedia, MM '18*, pages 1337–1345, New York, NY, USA, 2018. ACM. event-place: Seoul, Republic of Korea. 1
- [37] K. Simonyan and A. Zisserman. Two-Stream Convolutional Networks for Action Recognition in Videos. In *NIPS*, pages 568–576. 2014. 2
- [38] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, pages 2818–2826, 2016. 3
- [39] J. Thies, M. Zollhofer, M. Stamminger, C. Theobalt, and M. Niessner. Face2face: Real-Time Face Capture and Reenactment of RGB Videos. In *CVPR*, pages 2387–2395, 2016. 2, 3
- [40] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, pages 4489–4497, 2015. 2
- [41] S. Vosoughi, D. Roy, and S. Aral. The spread of true and false news online. *Science*, 359(6380):1146–1151, 2018. 1
- [42] W. Wang, Z. Cui, Y. Yan, J. Feng, S. Yan, X. Shu, and N. Sebe. Recurrent face aging. In *CVPR*, pages 2378–2386, 2016. 2
- [43] W. Wu, M. Kan, X. Liu, Y. Yang, S. Shan, and X. Chen. Recursive spatial transformer (rest) for alignment-free face recognition. In *ICCV*, pages 3772–3780, 2017. 3
- [44] Y. Wu, W. Abd-Almageed, and P. Natarajan. Deep Matching and Validation Network: An End-to-End Solution to Constrained Image Splicing Localization and Detection. In *ACMMM*, pages 1480–1502, New York, NY, USA, 2017. 1
- [45] Y. Wu, W. Abd-Almageed, and P. Natarajan. BusterNet: Detecting Copy-Move Image Forgery with Source/Target Localization. In *ECCV*, pages 168–184, 2018. 1, 3
- [46] Y. Wu, W. Abd-Almageed, and P. Natarajan. Image copy-move forgery detection via an end-to-end deep neural network. In *WACV*, pages 1907–1915. IEEE, 2018. 3
- [47] Y. Wu, W. AbdAlmageed, and P. Natarajan. ManTra-Net: Manipulation tracing network for detection and localization of image forgeries with anomalous features. In *CVPR*, June 2019. 3
- [48] F. Yu, V. Koltun, and T. Funkhouser. Dilated residual networks. In *CVPR*, pages 472–480, 2017. 3
- [49] P. Zhou, X. Han, V. I. Morariu, and L. S. Davis. Two-stream neural networks for tampered face detection. In *CVPR Workshops*, pages 1831–1839. IEEE, 2017. 1, 2