

Detecting Audio-Visual Deepfakes with Fine-Grained Inconsistencies

Marcella Astrid¹
marcella.astrid@uni.lu

Enjie Ghorbel^{1,2}
enjie.ghorbel@isamm.uma.tn

Djamila Aouada¹
djamila.aouada@uni.lu

¹ Computer Vision, Imaging & Machine Intelligence Research Group (CVI²), Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg, Luxembourg

² Cristal Laboratory, National School of Computer Sciences, Manouba University, Tunisia

Abstract

Existing methods on audio-visual deepfake detection mainly focus on high-level features for modeling inconsistencies between audio and visual data. As a result, these approaches usually overlook finer audio-visual artifacts, which are inherent to deepfakes. Herein, we propose the introduction of fine-grained mechanisms for detecting subtle artifacts in both spatial and temporal domains. First, we introduce a local audio-visual model capable of capturing small spatial regions that are prone to inconsistencies with audio. For that purpose, a fine-grained mechanism based on a spatially-local distance coupled with an attention module is adopted. Second, we introduce a temporally-local pseudo-fake augmentation to include samples incorporating subtle temporal inconsistencies in our training set. Experiments on the DFDC and the FakeAVCeleb datasets demonstrate the superiority of the proposed method in terms of generalization as compared to the state-of-the-art under both in-dataset and cross-dataset settings.

1 Introduction

In recent years, the capabilities of generative techniques, especially deep learning-based methods, in creating audio-visual deepfake data have rapidly improved [19, 27, 65]. Despite their advantages, these techniques can also be damaging to society if used with malicious intent. For example, a finance worker was recently scammed out of 25 million US dollars after engaging in a video call with a deepfake of the company's chief financial officer [4]. Therefore, it is essential to propose detection methods that are capable of identifying deepfakes.

To detect audio-visual deepfakes, one can exploit the inconsistencies between audio and visual data present in deepfakes. Several methods have been developed by focusing on this aspect [8, 12, 13]. Nevertheless, despite their effectiveness, these methods have not explored




Figure 1: (a) Previous works [8, 13] utilize high-level global features to measure inconsistencies between audio and visual data. (b) The proposed method measures the inconsistency between different visual regions and the audio input.




Figure 2: (a) The proposed temporally-local pseudo-fake synthesis involves the replacement of a small video segment by a subsequence extracted from another video (marked in blue). (b) The same strategy is followed for audio data.

fine-grained audio-visual representations. In fact, they only operate on high-level global features, while deepfake artifacts are known to be typically localized [34, 44]. In this work, we posit that by employing fine-grained strategies for modeling subtle audio-visual artifacts, deepfakes can be detected more effectively.

Based on this assumption, we introduce novel audio-visual fine-grained mechanisms tailored to deepfake detection at both the spatial and the temporal levels. To the best of our knowledge, no work has explored fine-grained inconsistencies to detect audio-visual deepfakes. In the spatial domain, instead of measuring inconsistencies through high-level representations as done in previous works [8, 13] (see Figure 1(a)), we consider features extracted from different spatial patches (Figure 1(b)). Attention is also incorporated, as we believe that only specific regions of the spatial domain are relevant to deepfake detection. In the temporal domain, we propose augmenting the fake data by simulating audio-visual inconsistencies. However, instead of applying augmentations to the entire audio-visual stream as in [12], we only manipulate a small portion of the sequence to simulate subtle artifacts, as illustrated in Figure 2. We conduct experiments on two audio-visual deepfake detection benchmarks, namely, DFDC and FakeAVCeleb. The results demonstrate enhanced generalization capabilities as compared to the state-of-the-art under both in-dataset and cross-dataset settings.

In summary, our contributions are listed below: 1) We propose a spatially-local architecture for detecting audio-visual deepfakes by implicitly measuring inconsistencies between sub-regions of the visual data and the audio; 2) We apply an additional cross-attention mechanism that enforces the model to focus on the inconsistency-prone visual regions; 3) We propose a temporally-local pseudo-fake augmentation that simulates fine-grained audio-visual artifacts; 4) We train our model on the DFDC dataset [11] and evaluate it for both in-dataset (DFDC) and cross-dataset (FakeAVCeleb dataset [23]) settings, demonstrating enhanced generalization capabilities in comparison to state-of-the-art (SoA) methods.

Paper organization: Section 2 discusses related work and positions our work with respect to the existing literature. In Section 3, we present a detailed description of the proposed method. Section 4 outlines the experimental setup and presents the obtained results. Finally, Section 5 summarizes our findings and concludes this paper.

2 Related work

2.1 Audio-visual deepfake detection

There exist numerous approaches for detecting audio-visual deepfakes. One approach focuses on specific identities, targeting the detection of deepfakes related to certain individuals [2, 3, 9]. While useful for some applications, these methods are limited to detecting deepfakes of individuals included in the training set. In contrast, we aim for a more universal approach capable of detecting individual-agnostic deepfakes. Another approach integrates information from both audio and visual modalities [25, 26, 28, 30, 32, 40, 43, 45]. However, combining features may induce redundancy, which has been shown to compromise the generalization capabilities [34]. This is also confirmed by our experiments where the use of residual connections to inject redundant information has led to a decrease in performance. A third approach for multi-modal deepfake detection is to aggregate the predictions resulting from separate visual and audio models [18]. Then, if at least one model predicts the input as fake, the overall prediction is set to fake. However, this mechanism may lead to a large number of false positives. One method in deepfake detection utilizes audio and visual modalities to pretrain a model [15]. However, the final model is fine-tuned only on visual data and hence cannot be used for audio-visual deepfake detection.

Different from the aforementioned methods, our work is primarily related to approaches that exploit inconsistencies between audio and visual cues present in fake data. Existing methods have addressed this issue in numerous ways. Chugh *et al.* [8] aim to minimize the distance between audio and visual features extracted from real data while maximizing it otherwise. However, this matching primarily operates on high-level features, potentially overlooking subtle artifacts. Similarly, Gu *et al.* [13] restrict their analysis to the lips only, neglecting potential artifacts elsewhere on the face. Mittal *et al.* [33] focus on detecting mismatches in high-level emotions between visual and audio. Nevertheless, this approach heavily relies on emotion recognition models, which could be sub-optimal for deepfake detection. Feng *et al.* [12] employ an auto-regressive model to predict the synchronization of visual and audio pairs. Nonetheless, their approach involves the translation of the entire audio/visual input to mimic inconsistencies, potentially disregarding more localized discrepancies. In contrast, we propose a fine-grained method for detecting subtle audio-visual mismatches taking into account both the spatial and the temporal dimensions.

2.2 Pseudo-fake generation

Pseudo-fake generation has been widely acknowledged for its ability to enhance the generalization capabilities of deepfake detectors. In the visual domain, Li *et al.* [29] blend a face into another facial image. To reduce blending artifacts, Shiohara and Yamasaki [35] blend faces from the same individual, demonstrating significant improvements in generalization as compared to [29]. Mejri *et al.* [32] focus on blending specific facial regions, such as eyes or nose. Chen *et al.* [5] employ reinforcement learning to generate pseudo-fakes. In the temporal domain, Wang *et al.* [41] introduce temporal inconsistencies by dropping or repeating frames. In the audio-visual domain, Feng *et al.* [12] translate entire audios or image sequences to simulate inconsistencies. In this work, inspired by the findings of [35], we explore the generation of subtle pseudo-fake synthesis, but consider the simulation of audio-visual inconsistencies in the temporal domain.

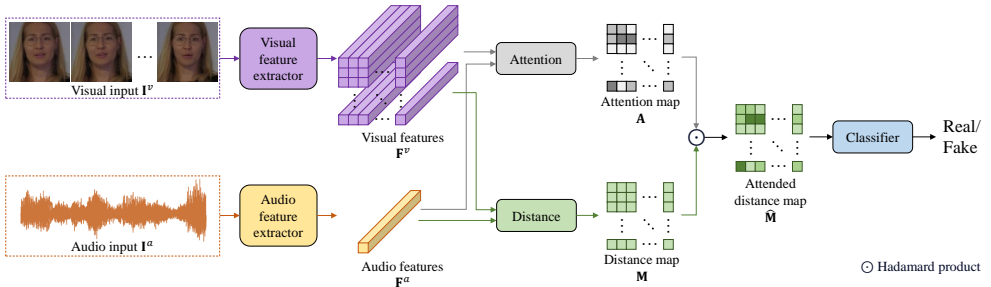


Figure 3: The proposed spatially-local deepfake detector: Firstly, audio and visual features are extracted, separately. Next, we compute the distance and attention maps between the audio and all spatial positions of the visual features. Subsequently, the distance map and the attention map are multiplied before being fed into a single-layer real/fake classifier.

2.3 Fine-grained deepfake detection

While simple binary deep neural networks have been initially employed for deepfake detection [56], they usually struggle to capture localized artifacts [42]. To overcome this limitation, several strategies have been proposed for visual deepfake detection. For instance, Zhao *et al.* [42] utilize multiple attention modules to capture fine-grained features across different regions of the input. On the other hand, Nguyen *et al.* [24] guide explicitly the network to focus on vulnerable points that are defined as the pixels which potentially suffer the most from blending artifacts. Nevertheless, to the best of our knowledge, no prior investigation has focused on examining inconsistencies between audio and local regions of the visual data for detecting audio-visual deepfakes.

3 Methodology

In order to focus on subtle artifacts, we propose two components: a spatially-local architecture for audio-visual deepfake detection (Section 3.1) and a temporally-local pseudo-fake data generation (Section 3.2).

3.1 Spatially-local audio-visual deepfake detector architecture

Figure 3 illustrates the overall architecture of the proposed deepfake detector, which extracts the spatially-local inconsistencies map to classify whether an input pair is fake or not. It comprises the feature extractor, the distance calculation, the attention module, and the classifier.

3.1.1 Feature extractor

The input to our feature extractor is a pair of audio and visual inputs, denoted as \mathbf{I}^a and \mathbf{I}^v , respectively. The audio is represented as a waveform of size $T^a \times C^a = 48000 \times 1$, where T^a and C^a denote the temporal and channel dimensions, respectively. We opt for the waveform representation to mitigate dependencies on frequency conversion processes, which can potentially reduce the robustness of the detector [20, 89]. The visual input consists of a

sequence of images of size $T^v \times C^v \times H^v \times W^v = 30 \times 3 \times 224 \times 224$, where T^v , C^v , H^v , and W^v represent the temporal (number of frames), the channel, the height, and the width dimensions, respectively.

Each input is fed into a specialized feature extractor denoted as $\mathcal{A}(\cdot)$ and $\mathcal{V}(\cdot)$ as follows,

$$\mathbf{F}^a = \mathcal{A}(\mathbf{I}^a), \quad \mathbf{F}^v = \mathcal{V}(\mathbf{I}^v), \quad (1)$$

where \mathbf{F}^a represents the audio feature of size $T^{a'} \times C^{a'} = 128 \times 15$ and \mathbf{F}^v represents the visual feature of size $T^{v'} \times C^{v'} \times H^{v'} \times W^{v'} = 128 \times 15 \times 28 \times 28$.

The feature extractor $\mathcal{V}(\cdot)$ is based on ResNetConv3D [16] that we modify to be shallow in order to obtain high-resolution features. Moreover, recent works in visual deepfake detection also suggest that visual deepfake artifacts can be effectively captured with a shallower network [14, 15]. The branch $\mathcal{A}(\cdot)$ is then based on a Conv1D architecture. It is designed such that $T^{a'}$ and $C^{a'}$ are equal to $T^{v'}$ and $C^{v'}$, respectively, so that we can calculate the distance between audio and visual features in the next step.

3.1.2 Spatially-local distance map

To measure fine-grained inconsistencies between the audio and different spatial visual patches, we create a distance map $\mathbf{M} = (M_{i,j})_{1 \leq i \leq H^{v'}, 1 \leq j \leq W^{v'}}$, where each element $M_{i,j}$ is calculated as follows,

$$M_{i,j} = d(\mathbf{f}^a, \mathbf{f}_{i,j}^v), \quad (2)$$

where $d(\cdot)$ represents a distance function. In our experiments, $d(\cdot)$ corresponds to the L_2 distance. \mathbf{f}^a and $\mathbf{f}_{i,j}^v$ denote the flattened versions of \mathbf{F}^a and $\mathbf{F}_{i,j}^v$ (i.e., \mathbf{F}^v at position (i, j)), respectively. Both vectors \mathbf{f}^a and $\mathbf{f}_{i,j}^v$ have a size of $T^{v'} \cdot C^{v'}$. Currently, we use L_2 distance following the work in [8] which utilizes L_2 distance to measure the distance between global features of audio and visual data. Nevertheless, other distance functions can also be explored in the future.

3.1.3 Attention module

Some visual regions of \mathbf{M} , such as hair and background, might be irrelevant to the quantification of audio-visual inconsistencies. Therefore, we propose to utilize an attention map \mathbf{A} of the same size as \mathbf{M} to implicitly reduce the contributions of these regions in the final classification. This is achieved through the following element-wise multiplication,

$$\hat{\mathbf{M}} = \mathbf{M} \odot \mathbf{A}, \quad (3)$$

where \odot represents the Hadamard/element-wise product.

The attention map $\mathbf{A} = (A_{i,j})_{1 \leq i \leq H^{v'}, 1 \leq j \leq W^{v'}}$ is calculated using a cross-attention-like mechanism as described below,

$$A_{i,j} = s \left(\frac{\mathcal{E}^a(\mathbf{F}^a) \cdot (\mathcal{E}^v(\mathbf{F}^v))_{i,j}}{T^{v'} C^{v'}} \right), \quad (4)$$

where s represents the softmax activation function, \mathcal{E}^a is a Conv1D operation with $C^{v'}/4$ filters, \mathcal{E}^v is a Conv3D operation also with $C^{v'}/4$ filters, and \cdot denotes the dot product. Division by $T^{v'} C^{v'}$ is a normalization process.

3.1.4 Classifier

Finally, the estimated distance map $\hat{\mathbf{M}}$ serves as the input to the classifier \mathcal{C} as depicted in the following equation,

$$y = \mathcal{C}(\hat{\mathbf{M}}), \quad (5)$$

where y tends towards 1 if the input is classified as fake, indicating a likely presence of inconsistencies in M , and approaches 0 otherwise. The classifier \mathcal{C} consists of flattening $\hat{\mathbf{M}}$, followed by a linear layer with a sigmoid activation function. Thus, \mathcal{C} only use the information of local inconsistencies to predict whether an audio-visual input is fake. Training is performed using a binary cross-entropy loss.

3.2 Temporally-local pseudo-fake synthesis

In previous work [12], the entire sequence of either audio or visual data is replaced with an alternative content for generating pseudo-fakes. However, such a technique usually produces low-quality pseudo-fakes, thereby hindering the generalization capabilities of the network. To address this, we propose a data synthesis method that replaces only a local portion of the audio, the visual, or both, to create multi-modal pseudo-fakes incorporating more subtle inconsistencies. This approach is inspired by image deepfake detection techniques, where it has been demonstrated that the use of imperceptible pseudo-fakes enhances the generalization capabilities [53] of deepfake detectors. Note that our method remains inherently different, as it focuses on subtle inconsistencies between audio and visual data; rather than visual artifacts.

Given two visual or audio sequences of length n denoted as $\mathbf{I} = \{\mathbf{I}_1, \mathbf{I}_2, \mathbf{I}_3, \dots, \mathbf{I}_n\}$ and $\mathbf{J} = \{\mathbf{J}_1, \mathbf{J}_2, \mathbf{J}_3, \dots, \mathbf{J}_n\}$, we randomly select a segment of length l , which is randomly chosen within the range $[l_{min}, l_{max}]$, where $l_{min} \leq l_{max}$. The values of l_{min} and l_{max} are determined as follows,

$$l_{min} = r_{min} \cdot n, \quad l_{max} = r_{max} \cdot n, \quad (6)$$

where r_{min} and r_{max} are hyperparameters representing the length ratios, selected from the range $]0, 1]$. In addition to providing the locality extent, such a range also offers data diversity.

After fixing l , we randomly select a starting point $g \in [1, n - l]$. Subsequently, we replace the elements from \mathbf{I}_g to \mathbf{I}_{g+l} with the corresponding elements \mathbf{J}_g to \mathbf{J}_{g+l} . This replacement process can be applied to either visual data only, audio data only, or both. We incorporate the original data during training and dynamically synthesize pseudo-fakes with a probability of 0.5. The pseudo-fake data is considered as fake during the learning phase. An illustration of the local replacement process is provided in Figure 4.

4 Experiments

4.1 Experiment setup

Dataset. We evaluate our method using the DFDC [11] and the FakeAVCeleb [23] datasets. For DFDC, following the setup in [8, 53], we sample 15,300 training videos and 2,700 test videos. In the training set, we ensure a balanced distribution between real and fake videos by randomly selecting the same number of videos for each class, while in the test set, we maintain a distribution that is identical to the original dataset. The testing protocol where




Figure 4: Our temporally-local pseudo-fake data synthesis: Given the original dataset illustrated in (a), we can create three types of pseudo-fakes: modifying only the audio data, modifying only the visual data, or modifying both the audio and visual inputs, as illustrated in (b).

	Att.	PF	RC	DFDC	FAV
(a)				97.11%	56.83%
(b)		✓		94.47%	<u>71.24%</u>
(c)	✓			98.09%	58.57%
(d)	✓	✓		<u>97.81%</u>	82.51%
(e)	✓	✓	✓	93.45%	60.27%

Table 1: Ablation study of our work reported in terms of AUC on DFDC and FakeAVCeleb (FAV) datasets. "Att.", "PF", and "RC" represent Attention, Pseudo-Fakes, and Residual Connections, respectively. The results produced by **our method** are reported in (d). The best and second-best performances are marked with bold and underlined, respectively.




Figure 5: AUC values under different (a) r_{min} , (b) r_{max} , and (c) map size ($H^{v'}$ and $W^{v'}$) settings on the DFDC dataset (in-dataset) and the FakeAVCeleb dataset (cross-dataset). Here, $r_{min} \approx 0$ corresponds to $l_{min} = 2$. In the in-dataset setting, no significant difference is observed, while the opposite is observed in the cross-dataset setting.

we use the test split of DFDC is referred to as the *in-dataset* setup. For the *cross-dataset* setup, we utilize FakeAVCeleb solely during testing. We adopt the testing settings described in [22], where the test set comprises 70 real and 70 fake videos. We preprocess our dataset similarly to [8], except that we use the audio waveform format normalized to a range of -1 to 1 using min-max normalization. We train our model with DFDC, similar to the setup in [8, 33]. Nevertheless, training on FakeAVCeleb can also be explored in the future.

Feature extractor architecture. For the visual feature extractor $\mathcal{V}(\cdot)$, we modify the visual stream used in [8] by only using the first three convolution blocks. We build the audio feature extractor $\mathcal{A}(\cdot)$ using several Conv1D, batch norm, max pool, and ReLU layers.

Evaluation criteria. We assess the performance of our model using the video-level Area Under the ROC Curve (AUC). Video-level prediction is obtained by averaging the predictions of multiple input subsequences.

Parameters and implementation details. The model is trained using the Adam optimizer [24] for 100 epochs, with a learning rate of 10^{-3} , weight decay of 10^{-5} , and a batch size of 8. The model with the lowest loss across the 100 epochs is selected for testing. To ensure a wide variety of pseudo-fake samples, unless specified otherwise, we set r_{min} to a value close to zero ~ 0 (resulting in $l_{min} = 2$) and r_{max} to 1.

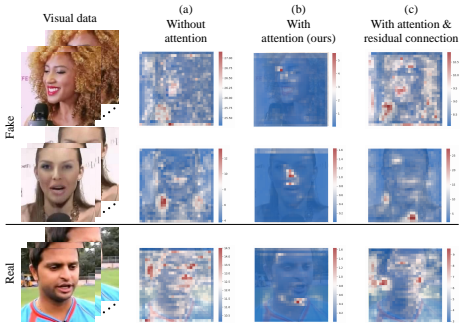


Figure 6: Visualization of $\hat{\mathbf{M}}$ extracted from three samples of the FakeAVCeleb dataset. Figures (a), (b), and (c) correspond to the same settings used in Table 1(b), (d), and (e), respectively. When the proposed attention mechanism is used without residual connections, the map is more focused on specific parts as compared to other settings.

4.2 Ablation study

4.2.1 Architecture design

For demonstrating the relevance of the proposed architecture, we conduct an ablation study and report the obtained results on DFDC and FakeAvCeleb with and without attention (i.e., $\hat{\mathbf{M}} = \mathbf{M}$) and using a residual connection (i.e., $\hat{\mathbf{M}} = (\mathbf{M} \odot \mathbf{A}) + \mathbf{M}$) to simulate the effect of introducing redundant information to the classifier.

Specifically, Table 1(b), (d), and (e) present the results of the model without attention, with attention, and with attention and a residual connection, respectively. It can be observed that the best results are obtained in Table 1(d) under both in-dataset and cross-dataset settings. For a deeper analysis of the results, we visualize the map $\hat{\mathbf{M}}$ in Figure 6. In Figure 6(b), we can observe that by only using attention, the model is able to better focus on specific regions as compared to the other setups. This suggests that attention allows disregarding some irrelevant parts such as the background; therefore leading to a more effective model. This gain in performance is also confirmed in Figure 7, where the distance between audio and visual features extracted from a real video tend to be lower when leveraging attention only, resulting in better discrimination between fake and real samples.

Moreover, for analyzing the impact of the proposed map distance, we also conducted experiments with different map sizes ($H^{v'}$ and $W^{v'}$, where $H^{v'} = W^{v'}$). To achieve that, we have applied an adaptive average pooling to \mathbf{F}^v before calculating the distance with the audio features, as applying global average pooling before classifier is a common practice in computer vision [17]. Figure 5(c) shows that even though there is generally no significant difference in performance when varying map sizes in the in-dataset setup, there is a significant drop in performance when the features become too global ($H^{v'} = W^{v'} = 1$) in the cross-dataset setting.

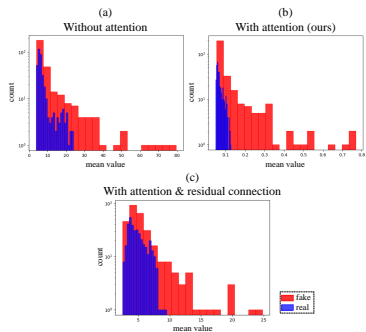


Figure 7: Histograms illustrating the distribution of the mean value of $\hat{\mathbf{M}}$ for real and fake data from the FakeAVCeleb test split. In Figures (a), (b), and (c), the same settings used in Table 1(b), (d), and (e) are considered, respectively. The separation between the distributions of real and fake data is more pronounced in Figure (b) as compared to Figures (a) and (c).

Method	Modality	FakeAVCeleb
LipForensics [14]	V	49.2%
CViT [14]	V	45.5%
MesoNet [14]	V	54.1%
MDS [14]	AV	72.9%
AVoiD-DF [14]	AV	82.8%
AVT ² -DWF [14]	AV	77.2%
Ours	AV	84.5%

Table 2: AUC under the cross-dataset setting, i.e., training on DFDC and testing on FakeAVCeleb. "V" and "AV" refer to visual-only and audio-visual modalities, respectively. The best performance is marked with bold.

Method	Modality	DFDC	Method	Modality	DFDC
Multi-attention [14]	V	84.8%	MDS [14]	AV	90.7%
SLADD [14]	V	75.2%	Emotion [14]	AV	84.4%
LipForensics [14]	V	73.5%	BA-TFD [14]	AV	84.6%
CViT [14]	V	63.7%	AVFakeNet [14]	AV	86.2%
MesoNet [14]	V	75.3%	VFD [14]	AV	85.1%
AASIST [14]	A	68.4%	AVoiD-DF [14]	AV	94.8%
ECAPA-TDNN [14]	A	69.8%	AVT ² -DWF [14]	AV	89.2%
RawNet [14]	A	56.2%	Ours	AV	97.7%

Table 3: AUC under the in-dataset setting, i.e., training on DFDC and testing also on DFDC. "V", "A", and "AV" refer to visual-only, audio-only, and audio-visual modalities, respectively. The best performance is marked with bold.



Figure 8: Visualization of the distance map \mathbf{M} , the attention map \mathbf{A} , and the attention-aware distance map $\hat{\mathbf{M}}$ for several examples from the FakeAVCeleb dataset. It is observed that \mathbf{A} reduces the impact of irrelevant parts in \mathbf{M} , resulting in a more focused attention through $\hat{\mathbf{M}}$.

4.2.2 Pseudo-fakes

In this subsection, we explore the relevance of the temporally-local pseudo-fake synthesis as well as the influence of r_{min} and r_{max} .

Specifically, we compare the results of the models without pseudo-fakes (Table 1(a) and (c)) and to the ones trained on pseudo-fakes (Table 1(b) and (d)). The substantial improvement observed in FakeAVCeleb suggests enhanced generalization capabilities when utilizing pseudo-fake data. While the performance on DFDC slightly decreases when incorporating pseudo-fakes, the drop in performance remains negligible as compared to the obtained improvement under the most relevant scenario i.e., the cross-dataset setting, especially when using attention module.

Figure 5(a) and (b) illustrate the impact of r_{min} and r_{max} on the performance, respectively. In the in-dataset scenario, no significant difference is observed when adjusting r_{min} and r_{max} . However, in the cross-dataset scenario, high values of r_{min} notably degrade the performance, emphasizing the importance of temporal locality achieved through the generation of pseudo-

fake data. Conversely, low values of r_{max} lead to a decreased performance, highlighting the importance of diversity in the pseudo-fake data.

4.3 Comparisons with state-of-the-art

Table 2 and Table 3 compare the proposed approach with SoA in terms of AUC on cross-dataset and in-dataset settings, respectively. For this comparison, we use the best performing model with $r_{min} = 0.3$ and $r_{max} = 1$. It can be seen that in general methods that are based on two modalities outperform single-modality methods. We also note that our method achieves better performance than audio-visual techniques, including, inconsistency-based methods that use more global features in the visual data, e.g., MDS [8] and Emotion [33].

4.4 Qualitative results

For a deeper understanding of the proposed method, we show in Figure 8 the distance map \mathbf{M} , attention map \mathbf{A} , and the attention-aware distance map $\hat{\mathbf{M}}$ extracted from several samples of the FakeAVCeleb dataset. As observed, \mathbf{M} may occasionally exhibit high distances in irrelevant parts, such as the background. However, \mathbf{A} reduces the impact of some irrelevant zones, allowing the consideration of more important portions. This trend is observed in both fake and real data. It can also be noted that the focus of the distance can be roughly similar between real and fake data. However, as observed in Figure 7(b), the overall distance map is different, creating a separation between real and fake data.

5 Conclusion

In this paper, a novel fined-grained method for audio-visual deep detection is proposed. Instead of measuring the inconsistency between global audio and visual features, more local strategies are considered. First, we propose a spatially-local architecture that computes the inconsistency between relevant visual patches and audio. Second, a temporally-local pseudo-fake data synthesis is introduced. The generated pseudo-fakes incorporating subtle inconsistencies are then used for training the proposed architecture. Experiments demonstrate the importance of the proposed components and their competitiveness as compared to the SoA. Several future directions include exploring different design choices and datasets.

Acknowledgment

This work was supported by the Luxembourg National Research Fund (FNR) under the project BRIDGES2021/IS/16353350/FaKeDeTeR and POST Luxembourg.

References

- [1] Darius Afchar, Vincent Nozick, Junichi Yamagishi, and Isao Echizen. Mesonet: a compact facial video forgery detection network. In *2018 IEEE international workshop on information forensics and security (WIFS)*, pages 1–7. IEEE, 2018.

- [2] Shruti Agarwal, Liwen Hu, Evonne Ng, Trevor Darrell, Hao Li, and Anna Rohrbach. Watch those words: Video falsification detection using word-conditioned facial motion. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 4710–4719, 2023.
- [3] Zhixi Cai, Kalin Stefanov, Abhinav Dhall, and Munawar Hayat. Do you really mean that? content driven audio-visual deepfake dataset and multimodal method for temporal forgery localization. In *2022 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, pages 1–10. IEEE, 2022.
- [4] Heather Chen and Kathleen Magramo. Finance worker pays out \$25 million after video call with deepfake ‘chief financial officer’. *CNN*. URL <https://edition.cnn.com/2024/02/04/asia/deepfake-cfo-scam-hong-kong-intl-hnk/index.html>.
- [5] Liang Chen, Yong Zhang, Yibing Song, Lingqiao Liu, and Jue Wang. Self-supervised learning of adversarial example: Towards good generalizations for deepfake detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18710–18719, 2022.
- [6] Harry Cheng, Yangyang Guo, Tianyi Wang, Qi Li, Xiaojun Chang, and Liqiang Nie. Voice-face homogeneity tells deepfake. *arXiv preprint arXiv:2203.02195*, 2022.
- [7] Harry Cheng, Yangyang Guo, Tianyi Wang, Qi Li, Xiaojun Chang, and Liqiang Nie. Voice-face homogeneity tells deepfake. *ACM Transactions on Multimedia Computing, Communications and Applications*, 2023.
- [8] Komal Chugh, Parul Gupta, Abhinav Dhall, and Ramanathan Subramanian. Not made for each other-audio-visual dissonance-based deepfake detection and localization. In *Proceedings of the 28th ACM international conference on multimedia*, pages 439–447, 2020.
- [9] Davide Cozzolino, Alessandro Pianese, Matthias Nießner, and Luisa Verdoliva. Audio-visual person-of-interest deepfake detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 943–952, 2023.
- [10] Brecht Desplanques, Jenthe Thienpondt, and Kris Demuynck. Ecapa-tdnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification. *arXiv preprint arXiv:2005.07143*, 2020.
- [11] Brian Dolhansky, Joanna Bitton, Ben Pflaum, Jikuo Lu, Russ Howes, Menglin Wang, and Cristian Canton Ferrer. The deepfake detection challenge (dfdc) dataset. *arXiv preprint arXiv:2006.07397*, 2020.
- [12] Chao Feng, Ziyang Chen, and Andrew Owens. Self-supervised video forensics by audio-visual anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10491–10503, 2023.
- [13] Yewei Gu, Xianfeng Zhao, Chen Gong, and Xiaowei Yi. Deepfake video detection using audio-visual consistency. In *Digital Forensics and Watermarking: 19th International Workshop, IWDW 2020, Melbourne, VIC, Australia, November 25–27, 2020, Revised Selected Papers 19*, pages 168–180. Springer, 2021.

- [14] Alexandros Haliassos, Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Lips don't lie: A generalisable and robust approach to face forgery detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5039–5049, 2021.
- [15] Alexandros Haliassos, Rodrigo Mira, Stavros Petridis, and Maja Pantic. Leveraging real talking faces via self-supervision for robust forgery detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14950–14962, 2022.
- [16] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6546–6555, 2018.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [18] Hafsa Ilyas, Ali Javed, and Khalid Mahmood Malik. Avfakenet: A unified end-to-end dense swin transformer deep learning model for audio–visual deepfakes detection. *Applied Soft Computing*, 136:110124, 2023.
- [19] Ye Jia, Yu Zhang, Ron Weiss, Quan Wang, Jonathan Shen, Fei Ren, Patrick Nguyen, Ruoming Pang, Ignacio Lopez Moreno, Yonghui Wu, et al. Transfer learning from speaker verification to multispeaker text-to-speech synthesis. *Advances in neural information processing systems*, 31, 2018.
- [20] Jee-weon Jung, Hee-Soo Heo, Ju-ho Kim, Hye-jin Shim, and Ha-Jin Yu. Rawnet: Advanced end-to-end deep neural network using raw waveforms for text-independent speaker verification. *arXiv preprint arXiv:1904.08104*, 2019.
- [21] Jee-weon Jung, Hee-Soo Heo, Hemlata Tak, Hye-jin Shim, Joon Son Chung, Bong-Jin Lee, Ha-Jin Yu, and Nicholas Evans. Aasist: Audio anti-spoofing using integrated spectro-temporal graph attention networks. In *ICASSP. IEEE*, 2022.
- [22] Hasam Khalid, Minha Kim, Shahroz Tariq, and Simon S Woo. Evaluation of an audio-video multimodal deepfake dataset using unimodal and multimodal detectors. In *Proceedings of the 1st workshop on synthetic multimedia-audiovisual deepfake generation and detection*, 2021.
- [23] Hasam Khalid, Shahroz Tariq, Minha Kim, and Simon S Woo. Fakeavceleb: A novel audio-video multimodal deepfake dataset. In *Proc. Conf. Neural Inf. Process. Syst. Datasets Benchmarks Track*, 2021.
- [24] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [25] Pavel Korshunov and Sébastien Marcel. Speaker inconsistency detection in tampered video. In *2018 26th European signal processing conference (EUSIPCO)*, pages 2375–2379. IEEE, 2018.

- [26] Pavel Korshunov, Michael Halstead, Diego Castan, Martin Graciarena, Mitchell McLaren, Brian Burns, Aaron Lawson, and Sebastien Marcel. Tampered speaker inconsistency detection with phonetically aware audio-visual features. In *International conference on machine learning*, number CONF, 2019.
- [27] Iryna Korshunova, Wenzhe Shi, Joni Dambre, and Lucas Theis. Fast face-swap using convolutional neural networks. In *Proceedings of the IEEE international conference on computer vision*, pages 3677–3685, 2017.
- [28] John K Lewis, Imad Eddine Toubal, Helen Chen, Vishal Sandesera, Michael Lomnitz, Zigfried Hampel-Arias, Calyam Prasad, and Kannappan Palaniappan. Deepfake video detection based on spatial, spectral, and temporal inconsistencies using multimodal deep learning. In *2020 IEEE Applied Imagery Pattern Recognition Workshop (AIPR)*, pages 1–9. IEEE, 2020.
- [29] Lingzhi Li, Jianmin Bao, Ting Zhang, Hao Yang, Dong Chen, Fang Wen, and Bain-ing Guo. Face x-ray for more general face forgery detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5001–5010, 2020.
- [30] Michael Lomnitz, Zigfried Hampel-Arias, Vishal Sandesara, and Simon Hu. Multi-modal approach for deepfake detection. In *2020 IEEE Applied Imagery Pattern Recognition Workshop (AIPR)*, pages 1–9. IEEE, 2020.
- [31] Nesryne Mejri, Konstantinos Papadopoulos, and Djamila Aouada. Leveraging high-frequency components for deepfake detection. In *2021 IEEE 23rd International Workshop on Multimedia Signal Processing (MMSP)*. IEEE, 2021.
- [32] Nesryne Mejri, Enjie Ghorbel, and Djamila Aouada. Untag: Learning generic features for unsupervised type-agnostic deepfake detection. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023.
- [33] Trisha Mittal, Uttaran Bhattacharya, Rohan Chandra, Aniket Bera, and Dinesh Manocha. Emotions don’t lie: An audio-visual deepfake detection method using affective cues. In *Proceedings of the 28th ACM international conference on multimedia*, pages 2823–2832, 2020.
- [34] Dat Nguyen, Nesryne Mejri, Inder Pal Singh, Polina Kuleshova, Marcella Astrid, Anis Kacem, Enjie Ghorbel, and Djamila Aouada. Laa-net: Localized artifact attention network for high-quality deepfakes detection. *arXiv preprint arXiv:2401.13856*, 2024.
- [35] Yuval Nirkin, Yosi Keller, and Tal Hassner. Fsgan: Subject agnostic face swapping and reenactment. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7184–7193, 2019.
- [36] Andreas Rossler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. Faceforensics++: Learning to detect manipulated facial images. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1–11, 2019.

- [37] Davide Salvi, Honggu Liu, Sara Mandelli, Paolo Bestagini, Wenbo Zhou, Weiming Zhang, and Stefano Tubaro. A robust approach to multimodal deepfake detection. *Journal of Imaging*, 9(6):122, 2023.
- [38] Kaede Shiohara and Toshihiko Yamasaki. Detecting deepfakes with self-blended images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18720–18729, 2022.
- [39] Hemlata Tak, Jose Patino, Massimiliano Todisco, Andreas Nautsch, Nicholas Evans, and Anthony Larcher. End-to-end anti-spoofing with rawnet2. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6369–6373. IEEE, 2021.
- [40] Rui Wang, Dengpan Ye, Long Tang, Yunming Zhang, and Jiacheng Deng. Avt2-dwf: Improving deepfake detection with audio-visual fusion and dynamic weighting strategies. *arXiv preprint arXiv:2403.14974*, 2024.
- [41] Zhendong Wang, Jianmin Bao, Wengang Zhou, Weilun Wang, and Houqiang Li. Altfreezing for more general video face forgery detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4129–4138, 2023.
- [42] Deressa Wodajo and Solomon Atnafu. Deepfake video detection using convolutional vision transformer. *arXiv preprint arXiv:2102.11126*, 2021.
- [43] Wenyuan Yang, Xiaoyu Zhou, Zhikai Chen, Bofei Guo, Zhongjie Ba, Zhihua Xia, Xiaochun Cao, and Kui Ren. Avoid-df: Audio-visual joint learning for detecting deepfake. *IEEE Transactions on Information Forensics and Security*, 18:2015–2029, 2023.
- [44] Hanqing Zhao, Wenbo Zhou, Dongdong Chen, Tianyi Wei, Weiming Zhang, and Nenghai Yu. Multi-attentional deepfake detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2185–2194, 2021.
- [45] Yipin Zhou and Ser-Nam Lim. Joint audio-visual deepfake detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14800–14809, 2021.