

AVFF: Audio-Visual Feature Fusion for Video Deepfake Detection

Trevine Oorloff^{1,2*} Surya Koppiseti² Nicolò Bonettini² Divyaraj Solanki²
Ben Colman² Yaser Yacoob¹ Ali Shahriyari² Gaurav Bharaj²

¹University of Maryland - College Park

²Reality Defender Inc.

Abstract

With the rapid growth in deepfake video content, we require improved and generalizable methods to detect them. Most existing detection methods either use uni-modal cues or rely on supervised training to capture the dissonance between the audio and visual modalities. While the former disregards the audio-visual correspondences entirely, the latter predominantly focuses on discerning audio-visual cues within the training corpus, thereby potentially overlooking correspondences that can help detect unseen deepfakes. We present Audio-Visual Feature Fusion (AVFF), a two-stage cross-modal learning method that explicitly captures the correspondence between the audio and visual modalities for improved deepfake detection. The first stage pursues representation learning via self-supervision on real videos to capture the intrinsic audio-visual correspondences. To extract rich cross-modal representations, we use contrastive learning and autoencoding objectives, and introduce a novel audio-visual complementary masking and feature fusion strategy. The learned representations are tuned in the second stage, where deepfake classification is pursued via supervised learning on both real and fake videos. Extensive experiments and analysis suggest that our novel representation learning paradigm is highly discriminative in nature. We report 98.6% accuracy and 99.1% AUC on the FakeAVCeleb dataset, outperforming the current audio-visual state-of-the-art by 14.9% and 9.9%, respectively.

1. Introduction

Deepfake generative AI technology enables new opportunities to create rich and quality content in multimedia applications such as virtual reality [52], movie production [41], and telepresence [35, 48]. However, its malicious use has become a major societal threat posing a number of problems including frauds¹, defamation² and disinformation³. As the

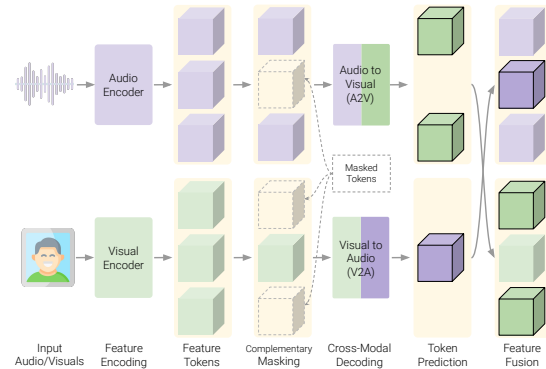


Figure 1. We use audio-visual correspondences for deepfake detection. Transformer-based encoders are used to extract audio and visual feature tokens, which are then masked complementarily. The visible audio tokens are sent through a learnable A2V network to predict the masked visual tokens. These predicted visual tokens are fused with the visible visual tokens to obtain the full visual embeddings. Full audio embeddings are obtained in a similar way using the V2A network. The audio/visual embeddings are then used for video reconstruction in the MAE sense, and subsequently for deepfake classification.

generative AI landscape continues to evolve, there is a growing need for robust deepfake detection that helps preserve content integrity. In this paper, we study video deepfake detection where either or both the visual and audio content are AI-generated.

We pursue multi-modal learning and draw inspiration from previous works, such as SyncNet [9], CLIP [42], and AudioCLIP [18], where the correspondence between different modalities (audio, text, visual) was leveraged to significantly enhance performance on various downstream tasks. We note that in real video face context, the audio-visual correspondence is deeply intuitive since there is an intrinsic correlation between the mouth articulations (visemes) and the speech units (phonemes) [2, 9, 41, 58], as well as an alignment of emotional nuances embedded in the facial and speech expressions [4, 36, 37]. Such inherent audio-visual correspondence, for example, in audio-driven emotion, is challenging to faithfully replicate in deepfake videos. Based on these observations, we propose a video deepfake detection method that learns efficient representations for audio and vi-

*This work was completed during an internship at Reality Defender Inc.

¹Fraudsters Used AI to Mimic CEO's Voice in Cybercrime Case

²Deepfake porn documentary explores its 'life-shattering' impact, AI Fake Nudes are booming

³Martin Lewis felt 'sick' seeing deepfake scam ad on Facebook, Deepfake scams have arrived

sual modalities. The proposed method employs a novel complementary masking and cross-modal feature fusion strategy to explicitly capture the audio-visual correspondences.

Previous literature on audio-visual video deepfake detection [36, 53] use supervised contrastive learning to capture the audio-visual correspondence. Such methods align the audio and visual embeddings closer to each other, if the content in both modalities is real, and push them apart if either or both modalities are generative. Similarly, others pursue a single stage supervised learning method, where models are directly trained on labeled deepfake datasets for deepfake classification [11, 29]. While such methods yield promising results, we conjecture that they may not fully exploit the audio-visual correspondence. Also, training solely on a deepfake dataset narrows the model’s focus to discern separable features within the training corpus, potentially overlooking subtle audio-visual correspondences that can help detect unseen deepfake samples (observe in Tabs. 1 and 3, the weaker performance of other baselines compared to [20] and ours).

To circumvent these issues, we propose a two stage training pipeline comprising of (i) a self-supervised representation learning stage that explicitly enforces audio-visual correspondence using a novel approach, and (ii) a supervised downstream classification stage. In the representation learning stage, we extract audio-visual representations via self-supervised learning on real face videos, which are available in abundance [1, 51, 59]. Drawing inspirations from CAV-MAE [17], we make use of the complementary nature of two learning objectives: contrastive learning and autoencoding. For extracting rich representations, we supplement the contrastive learning objective by a novel audio-visual complementary masking and fusion strategy that sits within the autoencoding objective. In the classification stage, we train a classifier that exploits the lack of cohesion between audio-visual features of deepfake videos to separate them from real videos.

We evaluate our method against existing state-of-the-art approaches on multiple benchmarks. Our results reveal substantial improvements, when compared against the existing audio-visual state-of-the-art, enhancing the performance by 9.9% in AUC and 14.9% in accuracy when evaluated on the FakeAVCeleb dataset [29]. This underscores the effectiveness of *explicitly* leveraging audio-visual correspondences through the proposed method. In summary:

- We propose a novel self-supervised representation learning method that *explicitly* captures audio-visual correspondences in real videos. To learn the correspondences, we pursue a dual-objective of contrastive learning and autoencoding, and supplement it with a novel audio-visual complementary masking and fusion strategy.
- Qualitative analysis using t-SNE [46] shows a clear separation between the real and fake video embeddings at the end of the representation learning stage. This demonstrates the

efficacy of the proposed representation learning.

- We propose a two-stage deepfake detection method comprising of the aforementioned representation learning stage followed by a deepfake classification stage. Our method yields state-of-the-art performance on deepfake detection when either or both the audio and visual contents are AI generated. We achieve 98.6% accuracy and 99.1% AUC on FakeAVCeleb, surpassing the existing audio-visual state-of-the-art by 14.9% and 9.9% respectively.

2. Related Works

2.1. Multi-Modal Representation Learning

Learning a joint representation from multiple modalities has been shown to be effective for different tasks in the state-of-the-art. SyncNet [9] proposes a Siamese Network to estimate the lip-sync error between audio and visual modalities. This framework processes each modality through a distinct branch and employs a contrastive loss to promote the similarities in the encoding space. More recently, improvements in the Natural Language Processing (NLP) field brought by BERT [10], allowed to use text modality in multi-modal frameworks. Another example is CLIP [42], a zero-shot image classification model that leverages separate encoders for images and captions to find a suitable pairing in the latent space. AudioCLIP [18] extends this approach to audio, enabling multi-modal classification.

Several self-supervised methods have emerged, inspired by the Masked Autoencoder (MAE) framework [21]. AV-MAE [16] is a joint masked autoencoder for audio, visual, and joint audio/visual classification. The authors explore different encoding policies for dual-modality inputs, demonstrating the ability to decode one masked modality from the other. CAV-MAE [17] raises concerns about the ability of a vanilla masked autoencoder to learn a coordinated representation between audio and visuals (*i.e.*, a representation that enforces similarity [5]) and adds a contrastive loss to explicitly leverage the audio-visual pair information. In this work, we draw inspiration from CAV-MAE, in using a dual contrastive-autoencoding objective for effective representation learning. Our approach diverges from existing MAE literature: (i) in terms of the masking strategy, where we use a complementary masking strategy post-encoding; (ii) in terms of the cross-modal fusion, where for every modality we replace the shared learnable masked tokens of MAEs with tokens predicted from the other modality. We do this to enforce explicit correspondence between audio and visual modalities.

2.2. Deepfake Detection

Visual-only methods. Multiple recent works have made use of visual-only artifacts for deepfake detection. The authors of [43] train a Convolutional Neural Network (CNN) (XceptionNet) end-to-end, setting one of the first baselines on the

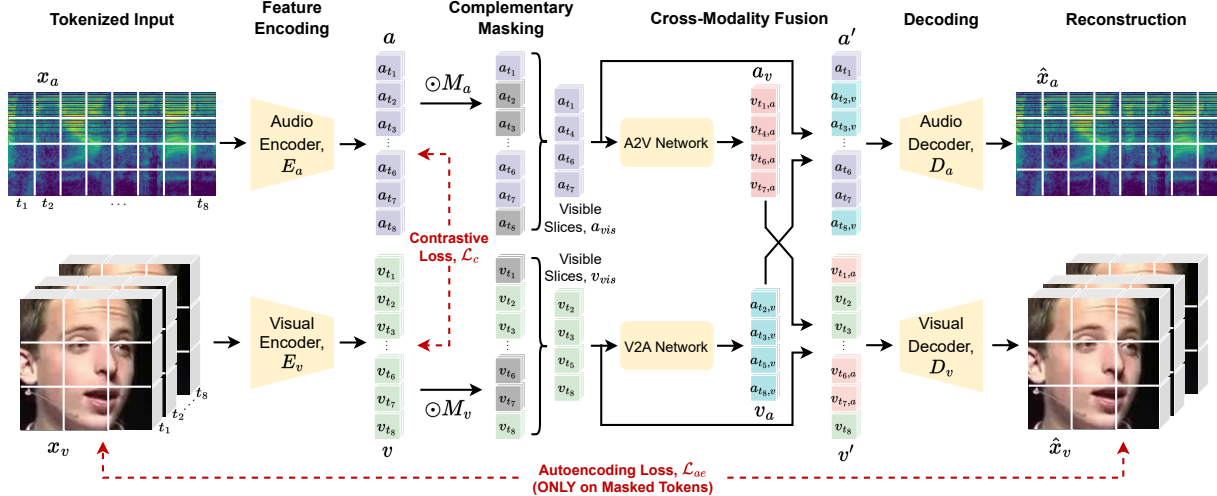


Figure 2. **Audio-Visual Representation Learning Stage.** A real input sample, $x \in \mathcal{D}_r$, with corresponding audio and visual tokens (x_a, x_v) , is split along the temporal dimension, creating K slices, $\{x_{a,t_i}\}_{i=1}^K$ and $\{x_{v,t_i}\}_{i=1}^K$ (illustrated with $K = 8$ in the figure). The temporal slices are then encoded using unimodal transformers, E_a and E_v , to yield feature embeddings a and v . We then complementarily mask 50% of the temporal slices in (a, v) with binary masks (M_a, M_v) . The visible slices of a and v are passed through A2V and V2A networks respectively, to generate cross-modal slices a_v and v_a . The masked slices of a and v are then replaced with the corresponding slices in a_v and v_a . The resulting cross-modality fusion representations, a' and v' , are input to unimodal decoders to obtain the audio and visual reconstructions, \hat{x}_a and \hat{x}_v . For the learning, we use a dual-objective loss function, which computes the contrastive loss between the audio and visual feature embeddings and the autoencoder loss between the input and the reconstruction of the masked tokens.

dataset they shared with the community. Some methods target specific face regions for exposing deepfakes. For example, LipForensics [19] relies on lip movements that might be difficult to reproduce by generative methods. Others consider inconsistent head pose [34, 54] or eye blinking [27, 33]. Another common approach is to consider both spatial and temporal domains. FTCN [56] proposes a combination of a CNN and a transformer network to exploit short-time and long-time temporal incoherence. Similarly, [55] extracts spatial features by means of an attention-based network and then fuses those features with a temporal module.

Several papers based on the Vision Transformer (ViT) have been published since the advent of the original paper [14] for image classification. An example is CViT [50], where learnable features are extracted by means of a CNN and subsequently fed to a ViT for the classification task. Similar approaches are followed by [13, 28, 60]. Recently, generalization to unseen deepfake methods [40, 49] and the impact of the identity leakage during training [12, 22] have also been investigated. RealForensics [20] proposed a hybrid approach that consists of using a multi-modal pre-training pipeline, where audio and visuals exclusively from real samples are used for computing internal representations that help the classifier to discriminate between real and fake video. This is not considered a pure multi-modal approach as the final classification is performed just on visuals, discarding the audio modality. As modern day video deepfakes consist of both audio and visual manipulations, uni-modal deepfake detection methods prove to be less effective.

Audio-visual methods. These methods consider audio and visual signals to target deepfake detection on both modalities. One of the first papers to address multi-modality is *Emotions Don't Lie* [36], proposing a Siamese Network where unimodal features are passed to an emotion recognition network to compare affective cues corresponding to perceived emotion from the two modalities within a video. *Not made for each other* [8] explicitly modeled the dissimilarity between modalities, and proposed the Modality Dissonance Score (MDS) network, where a contrastive loss is computed on single modality embeddings to expose differences on audio-visual pairs. *Voice-Face matching Detection* (VFD) [7] is another example of using contrastive loss for modeling face and voice homogeneity. A similar concept is exploited in [2], where the focus is on phoneme-viseme mismatch. The idea is that a given dynamic of mouth shape (viseme) should correspond to a given emitted sound (phoneme). The authors only focus on the mouth region, showing how deepfake methods struggle to reproduce certain dynamics.

More recently, the paradigm for multi-modality shifted towards fusion of single-modality features. AV-DFD [57] proposes a joint audio-visual deepfake detection framework in which visual and audio features are aligned and tiled to be passed onto a cross-attention mechanism on the temporal dimension. More recent papers study the encoding/decoding potential of ViTs and build feature fusion in the embedding space on the decoder side. Examples are AVFakeNet [24] and AVoid-DF [53].

3. Method

The proposed algorithm, AVFF, consists of two stages: (i) representation learning, and (ii) deepfake classification. Stage 1 aims to acquire an audio-visual representation with cross-modal correspondence via self-supervised learning, and it solely utilizes real face videos. The model learns audio-visual correspondences inherent to real videos via a contrastive learning objective and an effective complementary masking and fusion strategy that sits within an auto-encoding objective (see Fig. 2). Here, the complementary masking and fusion strategy takes uni-modal audio and visual embeddings (\mathbf{a}, \mathbf{v}) and systematically masks them to force the learning of advanced embeddings $(\mathbf{a}', \mathbf{v}')$ via reconstruction in the MAE sense. To instill cross-modal dependency, tokens from one modality are used to learn the masked embeddings of the other modality via cross-modal token conversion networks (A2V and V2A in Fig. 2). Since we work exclusively with real face videos in this stage, the model learns the dependency between “real” speech audio and the corresponding visual facial features. In Stage 2, a classifier is trained to distinguish between real and fake videos using the learned representations from the first stage. Effectively, the representation learning stage serves as pre-training for the downstream task of video deepfake detection.

3.1. Preprocessing

Initially, visual frames and the corresponding audio waveforms are extracted from raw videos at a sampling rate of 5 fps and 16 kHz, respectively. Given our emphasis on audio-visual correspondence, we align the cropped facial regions and eliminate the background in the visual frames using FaceX-Zoo [47]. This step is performed since background variations typically exhibit minimal correspondence with speech audio. Simultaneously, the audio waveform is converted into a log-mel spectrogram with L frequency bins. Henceforth, we refer to the preprocessed visual frames and log-mel spectrograms as the visual and audio input representations, respectively.

3.2. Representation Learning Stage

The primary objective of this stage is to learn a representation that effectively captures audio-visual feature correspondences inherently present in real videos (and different thereof from the correspondences in fake videos). Drawing inspiration from CAV-MAE [17], we propose a dual self-supervised learning approach that incorporates contrastive learning and autoencoding objectives. While contrastive learning helps build cross-modal correlations [18, 38, 42, 44], we found in preliminary experiments that relying solely on it does not establish a strong correspondence between the audio and visual modalities. Therefore, we supplement it with an autoencoding objective and embed a complementary masking

and cross-modal fusion strategy into the autoencoding framework. This allows us to learn rich cross-modal representations that result in improved deepfake detection (see Tab. 4). In Fig. 2, we illustrate the overall pipeline for our representation learning stage and discuss its key components next.

Input Tokenization. Given a dataset \mathcal{D}_r of real talking human portrait videos, we denote a video sample $x \in \mathcal{D}_r$ with a time duration T as comprising of audio and visual components, $x_a \in \mathbb{R}^{T_a \times L}$ and $x_v \in \mathbb{R}^{T_v \times C \times H \times W}$, respectively. In x_a , the (T_a, L) denote the number of audio frames and mel-frequency bins. In x_v , the (T_v, H, W, C) denote the number of visual frames, height, width, and number of channels. We choose T_a and T_v such that $T_a \cdot n_a = T_v \cdot n_v = T$, where n_a and n_v are the sampling frequencies of the audio and visual sequences. We tokenize x_a using 16×16 non-overlapping 2D patches (similar to Audio-MAE [23]), and x_v using $2 \times 16 \times 16$ non-overlapping 3D spatio-temporal patches (similar to MARLIN [6]). The resulting representations are denoted as \mathbf{x}_a and \mathbf{x}_v . Subsequently, we segment each of the tokenized representations into 8 equal temporal slices, $\mathbf{x}_a = \{x_{a,t_i}\}_{i=1}^8$ and $\mathbf{x}_v = \{x_{v,t_i}\}_{i=1}^8$, where the number of slices was decided empirically. This slicing preserves the temporal correspondence of each slice between modalities, as x_{a,t_i} and x_{v,t_i} correspond to the same time interval.

Feature Encoding. The two uni-modal audio and visual encoders, E_a and E_v , encode the tokenized inputs, \mathbf{x}_a and \mathbf{x}_v , and output uni-modal features \mathbf{a} and \mathbf{v} respectively: $\mathbf{p} = \{p_{t_i}\}_{i=1}^8 = E_p(\mathbf{x}_p + \text{pos}_p^e)$, where $p \in \{a, v\}$ and pos_p^e is the learnable positional embedding.

Complementary Masking. Within the uni-modal feature embeddings, \mathbf{a} and \mathbf{v} , we mask 50% of the temporal slices using binary masks, $(\mathbf{M}_a, \mathbf{M}_v) \in \{0, 1\}$, such that \mathbf{M}_a and \mathbf{M}_v are complementary to each other, *i.e.*, $\mathbf{M}_a = \mathbf{1}$ for slices where $\mathbf{M}_v = \mathbf{0}$ and vice-versa. In other words, for every masked slice in the audio feature, the corresponding visual slice is visible and vice versa. Let us denote the visible temporal slices as $\mathbf{p}_{vis} = \mathbf{M}_p \odot \mathbf{p}$ and the masked temporal slices as $\mathbf{p}_{msk} = (\neg \mathbf{M}_p) \odot \mathbf{p}$, where $p \in \{a, v\}$, \odot denotes the Hadamard product and \neg is the NOT operator.

Cross-Modal Fusion. Next, the *visible* temporal slices \mathbf{a}_{vis} and \mathbf{v}_{vis} are sent through learnable audio-to-visual (A2V) and visual-to-audio (V2A) networks to create their cross-modal temporal counterparts, $\mathbf{v}_a = \text{A2V}(\mathbf{a}_{vis})$ and $\mathbf{a}_v = \text{V2A}(\mathbf{v}_{vis})$, respectively. Here, \mathbf{v}_a contains $\{v_{t_i,a} = \text{A2V}(a_{t_i}), \forall t_i \text{ where } a_{t_i} \in \mathbf{a}_{vis}\}$, and similarly \mathbf{a}_v . Each of the A2V/V2A networks is composed of a single-layer MLP to match the number of tokens of the other modality followed by a single transformer block. The audio embedding \mathbf{a}' is then created using cross-modal fusion, wherein, we take the original feature \mathbf{a} and replace each masked slice with the corresponding slice of the same temporal index in the cross-

modal vector \mathbf{a}_v given by the V2A network (see Fig. 2). The visual embedding \mathbf{v}' is similarly obtained from the original feature \mathbf{v} and the cross-modal feature \mathbf{v}_a . Effectively, this process replaces the masked temporal slices of each modality with cross-modal slices generated from the corresponding temporal slices in the other modality. This is feasible due to our complementary masking strategy, as the masked slices of one modality are the visible slices of the other modality.

Decoding. The uni-modal audio and visual decoders, G_a and G_v , take \mathbf{a}' and \mathbf{v}' as input to generate the audio and visual reconstructions, $\hat{\mathbf{x}}_a = G_a(\mathbf{a}' + \text{pos}_a^g)$ and $\hat{\mathbf{x}}_v = G_v(\mathbf{v}' + \text{pos}_v^g)$, where pos_a^g and pos_v^g are learnable positional embeddings for each modality. The decoders use a transformer-based architecture and are entrusted with the task of effectively utilizing the mix of uni-modal slices and cross-modal slices present in \mathbf{a}' and \mathbf{v}' to generate the reconstructions for the two modalities.

Loss Functions. We use a dual objective loss, comprising an audio-visual contrastive loss and an autoencoding loss. The bi-directional audio-visual contrastive loss is defined as:

$$\mathcal{L}_c = - \sum_{\substack{p,q \in \{a,v\}, \\ p \neq q}} \frac{1}{2N} \sum_{i=1}^N \log \left[\frac{\exp(\|\bar{\mathbf{p}}^{(i)}\|^T \|\bar{\mathbf{q}}^{(i)}\|/\tau)}{\sum_{j=1}^N \exp(\|\bar{\mathbf{p}}^{(i)}\|^T \|\bar{\mathbf{q}}^{(j)}\|/\tau)} \right] \quad (1)$$

where $\bar{\mathbf{p}}^{(i)}$ is the mean latent vector across the patch dimension of the uni-modal embeddings of the i -th data sample, N is the number of samples, τ is the temperature parameter, and i, j are sample indices. The audio-visual contrastive loss enforces similarity constraints between the audio and visual embeddings of a given sample.

The autoencoder loss, \mathcal{L}_{ae} , is composed of reconstruction and adversarial losses, similar to MARLIN [6]. The reconstruction MSE loss, \mathcal{L}_{rec} , is computed between the inputs $(\mathbf{x}_a, \mathbf{x}_v)$, and their reconstructions $(\hat{\mathbf{x}}_a, \hat{\mathbf{x}}_v)$, and is computed only over the masked tokens following the approach in MAEs [6, 17, 23]:

$$\mathcal{L}_{rec} = \sum_{p \in \{a,v\}} \frac{1}{N} \sum_{i=1}^N \|\mathbf{x}_{p_{msk}}^{(i)} - \hat{\mathbf{x}}_{p_{msk}}^{(i)}\| \quad (2)$$

For the adversarial loss, \mathcal{L}_{adv} , we use the Wasserstein GAN loss [3] to supplement the reconstruction loss by enhancing the features captured in the reconstructions of each modality. Similar to the reconstruction loss, the adversarial loss is computed only on the masked tokens:

$$\mathcal{L}_{adv}^{(G)} = - \sum_{p \in \{a,v\}} \frac{1}{N} \sum_{i=1}^N D_p(\hat{\mathbf{x}}_{p_{msk}}) \quad (3)$$

$$\mathcal{L}_{adv}^{(D)} = \sum_{p \in \{a,v\}} \frac{1}{N} \sum_{i=1}^N (D_p(\hat{\mathbf{x}}_{p_{msk}}) - D_p(\mathbf{x}_{p_{msk}})) \quad (4)$$

Here, D_p denotes the discriminator of each modality, and the $\mathcal{L}_{adv}^{(D)}$ and $\mathcal{L}_{adv}^{(G)}$ represents the adversarial loss during the generator and the discriminator training steps respectively.




Figure 3. **Deepfake Classification Stage.** Given a sample $x \in \mathcal{D}_{df}$, comprising of audio and visual inputs \mathbf{x}_a and \mathbf{x}_v , we obtain the unimodal features (\mathbf{a}, \mathbf{v}) and the cross-modal embeddings $(\mathbf{a}_v, \mathbf{v}_a)$. For each modality, the unimodal and cross-modal embeddings are concatenated to obtain $(\mathbf{f}_a, \mathbf{f}_v)$. A classifier network is then trained to take $(\mathbf{f}_a, \mathbf{f}_v)$ as input and predict if the input is real or fake.

The overall training loss for the generative training step is as follows, where the λ_* parameters represent the corresponding loss weights:

$$\mathcal{L}^{(G)} = \lambda_c \mathcal{L}_c + \lambda_{rec} \mathcal{L}_{rec} + \lambda_{adv} \mathcal{L}_{adv}^{(G)} \quad (5)$$

Computing the autoencoding loss objective on the masked temporal slices strictly enforces the decoder to learn from the other modality, as the input embeddings for the decoder at masked indices are obtained from the other modality. This novel strategy explicitly enforces audio-visual correspondence supplementing the contrastive loss objective.

3.3. Deepfake Classification Stage

The goal of this stage is to detect video deepfakes, where either or both audio and visual modalities have been faked. For this, we use the encoders and the cross-modal networks trained in the representation learning phase. We train a classifier to tell real videos and deepfakes apart using a supervised learning approach. The classification pipeline is depicted in Fig. 3. Since the learned representations have a high audio-visual correspondence for real videos, we expect the classifier to exploit the lack of audio-visual cohesion of synthesized samples in distinguishing between real and fake.

Input Tokenization. The process followed in input tokenization is identical to Stage 1 except for the dataset used. In this stage we draw samples from a labeled deepfake dataset (\mathcal{D}_{df}) consisting of both real and fake videos, *i.e.*, $(x, y) \in \mathcal{D}_{df}$, where x is the video sample and y is the label (real/fake).

Feature Extraction. The tokenized inputs $(\mathbf{x}_a, \mathbf{x}_v)$, are sent through the backbone of Stage 1 to obtain (i) the feature embeddings (\mathbf{a}, \mathbf{v}) , which are the outputs of the uni-modal encoders, and (ii) the cross-modal embeddings $(\mathbf{a}_v, \mathbf{v}_a)$, which are the outputs of the A2V/V2A cross-modal networks. Here, the cross-modal embeddings are computed for all temporal slices (note: we do not use masking in this stage). We concatenate the two embeddings of each modality creating

$(\mathbf{f}_a, \mathbf{f}_v)$, where $\mathbf{f}_p = \mathbf{p} \oplus \mathbf{p}_q$, $\forall p, q \in \{a, v\}, p \neq q$, and \oplus is the concatenation operator along the feature dimension.

Classifier Network. The classifier network, Q , takes as input the combined embeddings of each modality $(\mathbf{f}_a, \mathbf{f}_v)$, and predicts if a given sample is real or fake. The classifier network consists of two uni-modal patch reduction networks: (Ψ_a, Ψ_v) , followed by a classifier head, Γ . Each embedding $(\mathbf{f}_a, \mathbf{f}_v)$, is first distilled in the patch dimension using the corresponding uni-modal patch reduction networks. Then the output embeddings are concatenated along the feature dimension and fed into the classifier head which outputs the logits, l , used to classify if a given sample is real or fake. Formally, $l = Q(\mathbf{f}_a, \mathbf{f}_v) = \Gamma(\Psi_a(\mathbf{f}_a) \oplus \Psi_v(\mathbf{f}_v))$.

Loss Function. We use the standard cross-entropy loss, denoted by \mathcal{L}_{CE} as the learning objective, computed using the input labels, y , and the output logits, l .

Deepfake Classifier Inference Stage. During inference, we first split the video into blocks of time T (the sample length during training) with a step size of $T/8$, which is the duration of a temporal slice. The output logits are computed for each of the blocks and the classification decision (real or fake) is made based on the mean of the output logits.

4. Experiments and Results

4.1. Implementation

We train Stage 1 (representation learning), using the LRS3 dataset [1], which exclusively contains real videos. In Stage 2 (deepfake classification), we train a classifier that follows a supervised learning approach using the FakeAVCeleb [29] dataset. FakeAVCeleb comprises of both real and fake videos, where either one or both audio-visual modalities have been synthesized using different combinations of several generative deepfake algorithms (visual: FaceSwap[31], FSGAN[39], and Wav2Lip [41]; audio: SV2TTS [26]). Please refer to the supplementary for additional details on datasets, architecture, and implementation.

4.2. Evaluation and Discussion

We evaluate the performance of our model against existing state-of-the-art algorithms, on multiple criteria: intra-dataset performance, cross-manipulation generalization, and cross-dataset generalization following [15, 53]. We compare our results against state-of-the-art audio-visual approaches and uni-modal (visual) approaches for completeness. We report accuracy (ACC), average precision (AP), and area under the ROC curve (AUC) averaged across multiple runs with different random seeds. For audio-visual algorithms, we label a video as fake if either or both audio and visual modalities have been manipulated. To maintain fairness, for uni-modal algorithms we consider a video as fake only if the visual modality has been manipulated. Please refer to the

Method	Modality	ACC	AUC
Xception [43]	V	67.9	70.5
LipForensics [19]	V	80.1	82.4
FTCN [56]	V	64.9	84.0
CViT [50]	V	69.7	71.8
RealForensics [20]	V	89.9	94.6
Emotions Don't Lie [36]	AV	78.1	79.8
MDS [8]	AV	82.8	86.5
AVFakeNet [24]	AV	78.4	83.4
VFD [7]	AV	81.5	86.1
AVoiD-DF [53]	AV	83.7	89.2
AVFF (Ours)	AV	98.6	99.1

Table 1. **Intra-Dataset Performance.** We evaluate our method against baselines using a 70%-30% train-test split on the FakeAVCeleb dataset, where we achieve state-of-the-art performance by significant margins. Best result is in bold, and second best per modality is underlined.

supplementary for additional results on robustness to unseen audio and visual perturbations.

Intra-Dataset Performance. Following the methodology outlined in [53], our training utilizes 70% of all FakeAVCeleb samples, while the remaining 30% constitutes the unseen test set. As denoted in Tab. 1, our approach demonstrates substantial improvements over the existing state-of-the-art, both in audio-visual (AVoiD-DF [53]) and uni-modal (RealForensics [20]) deepfake detection. Compared to AVoiD-DF, our method achieves an increase in accuracy of 14.9% (+9.9% in AUC), and compared to RealForensics the accuracy increases by 8.7% (+4.5% AUC). Overall, the superior performance of audio-visual methods leveraging cross-modal correspondence is evident, outperforming uni-modal approaches that rely on uni-modal artifacts (*i.e.* visual anomalies) introduced by deepfake algorithms. RealForensics, while competitive, discards the audio modality during detection, limiting its applicability exclusively to visual deepfakes. This hinders its practicality as contemporary deepfakes often involve manipulations in both audio and visual modalities. The enhanced results of both RealForensics and our proposed method highlight the positive impact of employing a pre-training stage for effective representation learning. This observation aligns with findings in other multi-modal representation learning research across diverse downstream tasks [6, 17, 23].

Cross-Manipulation Generalization. In this experiment, we aim to assess the model's performance on samples generated using previously unseen manipulation methods. The scalability of deepfake detection algorithms to unseen manipulation methods is crucial for adapting to evolving threats, thus ensuring wide applicability across diverse scenarios.

Similar to [15], we partition the FakeAVCeleb dataset into five categories: RVFA, FVRA-WL, FVFA-FS, FVFA-GAN,

Method	Modality	RVFA		FVRA-WL		FVFA-FS		FVFA-GAN		FVFA-WL		AVG-FV	
		AP	AUC	AP	AUC	AP	AUC	AP	AUC	AP	AUC	AP	AUC
Xception [43]	V	-	-	88.2	88.3	92.3	93.5	67.6	68.5	91.0	91.0	84.8	85.3
LipForensics [19]	V	-	-	97.8	97.7	99.9	99.9	61.5	68.1	98.6	98.7	89.4	91.1
FTCN [56]	V	-	-	96.2	97.4	100.	100.	77.4	78.3	95.6	96.5	92.3	93.1
RealForensics [20]	V	-	-	88.8	93.0	99.3	99.1	<u>99.8</u>	<u>99.8</u>	93.4	96.7	<u>95.3</u>	<u>97.1</u>
AV-DFD [57]	AV	74.9	73.3	<u>97.0</u>	97.4	99.6	99.7	58.4	55.4	100.	100.	88.8	88.1
AVAD (LRS2) [15]	AV	62.4	71.6	93.6	93.7	95.3	95.8	94.1	94.3	93.8	94.1	94.2	94.5
AVAD (LRS3) [15]	AV	70.7	<u>80.5</u>	91.1	93.0	91.0	92.3	91.6	92.7	91.4	93.1	91.3	92.8
AVFF (Ours)	AV	93.3	92.4	94.8	98.2	100.	100.	99.9	100.	<u>99.4</u>	<u>99.8</u>	98.5	99.5

Table 2. **Cross-Manipulation Generalization on FakeAVCeleb.** We evaluate the model’s performance by leaving out one category for testing while training on the rest. We consider the following 5 categories in FakeAVCeleb: (i) **RVFA**: Real Visual - Fake Audio (SV2TTS), (ii) **FVRA-WL**: Fake Visual - Real Audio (Wav2Lip), (iii) **FVFA-FS**: Fake Visual - Fake Audio (FaceSwap + Wav2Lip + SV2TTS), (iv) **FVFA-GAN**: Fake Visual - Fake Audio (FaceSwapGAN + Wav2Lip + SV2TTS), and (v) **FVFA-WL**: Fake Visual - Fake Audio (Wav2Lip + SV2TTS). The column titles correspond to the category of the test set. AVG-FV corresponds to the average metrics of the categories containing fake visuals. Best result is in bold, and second best is underlined. Our method yields consistently high performance across all manipulation methods while yielding state-of-the-art performance for several categories.

Method	Modality	AP	AUC
Xception [43]	V	76.9	77.7
LipForensics [19]	V	89.5	86.6
FTCN [56]	V	66.8	68.1
RealForensics [20]	V	95.7	<u>93.6</u>
AV-DFD [57]	AV	79.6	82.1
AVAD [15]	AV	87.6	86.9
AVFF (Ours)	AV	<u>93.1</u>	95.5

Table 3. **Cross-Dataset Generalization on KoDF.** We evaluate our model’s performance against baselines by testing the model trained on the FakeAVCeleb dataset, on a subset of the KoDF dataset. Best result is in bold, and second best is underlined.

and FVFA-WL, based on the algorithms used to generate the deepfakes. Descriptions of these categories are included in the caption of Tab. 2, for ease of reference. Using these categories, we evaluate the model leaving one category out for testing while training on the remaining categories. Results are reported in Tab. 2. Our method achieves the best performance in almost all cases (and at par with the rest) and notably, yields consistently enhanced performance (AUC > 92+%, AP > 93+%) across all categories, while other baselines (Xception [43], LipForensics [19], FTCN [56], AV-DFD [57]) fall short in categories FVFA-GAN and RVFA. While the performance of AVAD (an unsupervised method) is intriguing, we can notice how other supervised baselines perform better in most cases.

Cross-Dataset Generalization. We also evaluate the adaptability of the model to a different data distribution, by testing on a subset of the KoDF dataset [32], following the evaluation protocol established by [15]. We report the results in Tab. 3, where we perform at par with RealForensics. Overall, we observe a performance trend similar to cross-

generalization to unseen generative methods. Please refer to the supplementary for additional cross-dataset generalization results on DF-TIMIT [30] and DFDC [11] datasets.

Analysis on the Learned Representation. During the training of the downstream task, we observe notable AUC scores on the test set, reaching as high as 90% within the initial 1-3 epochs. Intrigued by this observation, we explore the representations learned at the end of Stage 1 (Sec. 3.2). We visualize the t-SNE plots of embeddings for random samples from each category of the FakeAVCeleb dataset in Fig. 4. We do not expose Stage 1 to any deepfake samples during training, and still observe clear discrimination between real and fake samples. This explains the high AUC scores achieved even at the very early stages of the downstream training. A consequence of our representation learning stage is that we observe a clear disentanglement between real and deepfake videos within the FakeAVCeleb dataset. Further, distinct clusters are evident for each deepfake category which indicates that our representations are capable of capturing subtle cues that differentiate different deepfake algorithms despite not encountering any of them during the training stage.⁴ A further analysis of the t-SNE visualizations reveals that the samples belonging to adjacent clusters are related in terms of the deepfake algorithms used to generate them. For instance, FVRA-WL and FVFA-WL, which are adjacent, both employ Wav2Lip to synthesize the deepfakes (refer to the encircled regions in Fig. 4). These findings underscore the efficacy of our novel audio-visual representation learning paradigm. Please refer to the supplementary for the evaluation of the classification performance on the learned representation which further reinforces the above analysis.

⁴Stage 1 is exclusively trained on the LRS3 dataset containing only real videos, while samples for the t-SNE are drawn from FakeAVCeleb.




Figure 4. **The t-SNE Visualization of the Embeddings at the end of the Representation Learning Stage.** A clear distinction is seen between the representations of real and fake videos, as well as between different deepfake categories. Further analysis indicates that samples of adjacent clusters are generated using the same deepfake algorithm, which we encircle manually to highlight the clusters.

5. Ablation Study

Autoencoding Objective. In this experiment, we train the model using only the contrastive loss objective, discarding the autoencoding objective. This deactivates complementary masking, cross-modality fusion, and decoding modules. The feature embeddings at the output of the encoders (\mathbf{a} , \mathbf{v}), are used for the downstream training. Results (see row (i) in Tab. 4) indicate a performance reduction, highlighting the importance of the autoencoding objective.

Cross-Modal Fusion. In this ablation, we discard the A2V/V2A networks, which predict the masked tokens of the other modality, and instead use shared learnable masked tokens similar to MAE approaches [6, 17, 21, 23]. We can see that the performance of the model diminishes (especially AP) (see row (ii) in Tab. 4). This highlights the importance of the cross-modal fusion module, as it supplements the representation of a given modality with information extracted from the other modality, helping building the correspondence between the two modalities.

Complementary Masking. Replacing complementary masking with random masking in Stage 1, results in a notable drop in AP and AUC scores, affecting the model’s ability to learn correspondences (see row (iii) in Tab. 4). We attribute this performance drop to the inability of the model to learn correspondences between audio and visual modalities due to the randomness, which indicates the importance of complementary masking in the proposed method.

Concatenation of Different Embeddings. In the deepfake classification stage, we concatenate the feature embeddings (\mathbf{a} , \mathbf{v}), with the cross-modal embeddings (\mathbf{a}_v , \mathbf{v}_a), creating the concatenated embeddings (\mathbf{f}_a , \mathbf{f}_v). In this experiment, we evaluate the performance using each of the embeddings in isolation (see rows (iv) and (v) in Tab. 4). While the use of each embedding generates promising results, the synergy of the two embeddings enhances the performance.

Method	AP	AUC
(i) Only contrastive loss	84.2	90.3
(ii) Ours w/o cross-modality fusion	87.2	93.1
(iii) Ours w/o complementary masking	78.9	90.7
(iv) Only feature embeddings	89.7	97.6
(v) Only cross-modal embeddings	94.6	98.0
(vi) Mean Pooling patch dimension	<u>96.5</u>	<u>98.1</u>
AVFF (Ours)	96.7	99.1

Table 4. **Evaluation on Ablations.** Best result is in bold, and second best is underlined. The proposed AVFF pipeline performs the best among the considered ablations.

Uni-Modal Patch Reduction. Replacing the uni-modal patch reduction networks (Ψ_a , Ψ_v) with Mean Pooling slightly dents the performance (see row (vi) in Tab. 4). This suggests that a simple linear averaging is sub-optimal in that some subtle cues may no longer be preserved when compared to using an MLP which effectively performs weighted non-linear averaging.

6. Conclusion, Limitations, and Future Work

In this paper, we propose AVFF, a novel two-staged learning framework for audio-visual deepfake detection. AVFF is composed of a self-supervised representation learning stage that captures the audio-visual correspondence using contrastive learning and a novel complementary masking and cross-modal fusion module within the autoencoding objective, followed by a supervised deepfake classification stage. Our method shows significant improvements in both in-distribution performance as well as generalization on unseen manipulations over both visual-only and audio-visual state-of-the-art algorithms. Our results not only validate the effectiveness of the proposed approach but also emphasize the potential as a defensive tool against the escalating threat posed by deepfake videos.

Limitations. Our work is limited to cases where there is a coherence between the audio and visual modalities, as we rely on it to distinguish fake from real videos. Hence, asynchronous videos with audio lags or mismatched audio, videos with multiple speakers, and visuals with occlusions (*e.g.*, masks, hands) would be challenging. Our model also requires the input to contain both audio and visual modalities, *i.e.*, videos with only one modality are not supported.

Future Works. A few possible future research avenues include: (i) supplementing our representation learning method with a strategy to preserve uni-modal cues to mitigate the aforementioned limitations, (ii) adaptation for other audio-visual downstream tasks such as emotion recognition, and (iii) generalization for non-humanoid audio-driven videos especially with diffusion-based approaches [25, 45].

Acknowledgements. Yaser Yacoob is supported in part by the DARPA SemaFor Program under HR001120C0124.

References

- [1] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. Lrs3-ted: a large-scale dataset for visual speech recognition. *arXiv preprint arXiv:1809.00496*, 2018. 2, 6
- [2] Shruti Agarwal, Hany Farid, Ohad Fried, and Maneesh Agrawala. Detecting deep-fake videos from phoneme-viseme mismatches. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 660–661, 2020. 1, 3
- [3] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR, 2017. 5
- [4] Weiming Bai, Yufan Liu, Zhipeng Zhang, Bing Li, and Weiming Hu. Aunet: Learning relations between action units for face forgery detection. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 24709–24719, 2023. 1
- [5] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2):423–443, 2019. 2
- [6] Zhixi Cai, Shreya Ghosh, Kalin Stefanov, Abhinav Dhall, Jianfei Cai, Hamid Rezatofighi, Reza Haffari, and Munawar Hayat. Marlin: Masked autoencoder for facial video representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1493–1504, 2023. 4, 5, 6, 8
- [7] Harry Cheng, Yangyang Guo, Tianyi Wang, Qi Li, Xiaojun Chang, and Liqiang Nie. Voice-face homogeneity tells deepfake. *arXiv preprint arXiv:2203.02195*, 2022. 3, 6
- [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. 3, 6
- [9] J. S. Chung and A. Zisserman. Out of time: automated lip sync in the wild. In *Workshop on Multi-view Lip-reading, ACCV*, 2016. 1, 2
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *North American Chapter of the Association for Computational Linguistics*, 2019. 2
- [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. 2, 7
- [12] Shichao Dong, Jin Wang, Renhe Ji, Jiajun Liang, Haoqiang Fan, and Zheng Ge. Implicit identity leakage: The stumbling block to improving deepfake detection generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3994–4004, 2023. 3
- [13] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Ting Zhang, Weiming Zhang, Nenghai Yu, Dong Chen, Fang Wen, and Baining Guo. Protecting celebrities from deepfake with identity consistency transformer. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9458–9468, 2022. 3
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020. 3
- [15] 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. 6, 7
- [16] Mariana-Iuliana Georgescu, Eduardo Fonseca, Radu Tudor Ionescu, Mario Lucic, Cordelia Schmid, and Anurag Arnab. Audiovisual masked autoencoders. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16144–16154, 2023. 2
- [17] Yuan Gong, Andrew Rouditchenko, Alexander H. Liu, David Harwath, Leonid Karlinsky, Hilde Kuehne, and James R. Glass. Contrastive audio-visual masked autoencoder. In *ICLR*. OpenReview.net, 2023. 2, 4, 5, 6, 8
- [18] Andrey Guzhov, Federico Raue, Jörn Hees, and Andreas Dengel. Audioclip: Extending clip to image, text and audio. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 976–980. IEEE, 2022. 1, 2, 4
- [19] 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. 3, 6, 7
- [20] 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. 2, 3, 6, 7
- [21] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15979–15988, 2022. 2, 8
- [22] Baojin Huang, Zhongyuan Wang, Jifan Yang, Jiabin Ai, Qin Zou, Qian Wang, and Dengpan Ye. Implicit identity driven deepfake face swapping detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4490–4499, 2023. 3
- [23] Po-Yao Huang, Hu Xu, Juncheng Li, Alexei Baevski, Michael Auli, Wojciech Galuba, Florian Metze, and Christoph Feichtenhofer. Masked autoencoders that listen. *Advances in Neural Information Processing Systems*, 35:28708–28720, 2022. 4, 5, 6, 8
- [24] Hafsa Ilyas, Ali Javed, and Khalid Mahmood Malik. Av-fakenet: A unified end-to-end dense swin transformer deep learning model for audio-visual deepfakes detection. *Applied Soft Computing*, 136:110124, 2023. 3, 6
- [25] Yujin Jeong, Wonjeong Ryoo, Seunghyun Lee, Dabin Seo, Wonmin Byeon, Sangpil Kim, and Jinkyu Kim. The power

- of sound (tpos): Audio reactive video generation with stable diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7822–7832, 2023. 8
- [26] 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. 6
- [27] Tackhyun Jung, Sangwon Kim, and Keecheon Kim. Deepvision: Deepfakes detection using human eye blinking pattern. *IEEE Access*, 8:83144–83154, 2020. 3
- [28] Bachir Kaddar, Sid Ahmed Fezza, Wassim Hamidouche, Zahid Akhtar, and Abdenour Hadid. Hcit: Deepfake video detection using a hybrid model of cnn features and vision transformer. In *2021 International Conference on Visual Communications and Image Processing (VCIP)*, pages 1–5, 2021. 3
- [29] Hasam Khalid, Shahroz Tariq, Minha Kim, and Simon S Woo. Fakeavceleb: A novel audio-video multimodal deepfake dataset. *arXiv preprint arXiv:2108.05080*, 2021. 2, 6
- [30] Pavel Korshunov and Sébastien Marcel. Deepfakes: a new threat to face recognition? assessment and detection. *arXiv preprint arXiv:1812.08685*, 2018. 7
- [31] 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. 6
- [32] Patrick Kwon, Jaeseong You, Gyuhyeon Nam, Sungwoo Park, and Gyeongsu Chae. Kodf: A large-scale korean deepfake detection dataset. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10744–10753, 2021. 7
- [33] Yuezun Li, Ming-Ching Chang, and Siwei Lyu. In ictu oculi: Exposing ai created fake videos by detecting eye blinking. In *2018 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–7, 2018. 3
- [34] Kevin Lutz and Robert Bassett. Deepfake detection with inconsistent head poses: Reproducibility and analysis. *CoRR*, abs/2108.12715, 2021. 3
- [35] Shugao Ma, Tomas Simon, Jason Saragih, Dawei Wang, Yuecheng Li, Fernando De la Torre, and Yaser Sheikh. Pixel codec avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 64–73, 2021. 1
- [36] 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. 1, 2, 3, 6
- [37] Trisha Mittal, Uttaran Bhattacharya, Rohan Chandra, Aniket Bera, and Dinesh Manocha. M3er: Multiplicative multimodal emotion recognition using facial, textual, and speech cues. In *Proceedings of the AAAI conference on artificial intelligence*, pages 1359–1367, 2020. 1
- [38] Pedro Morgado, Ishan Misra, and Nuno Vasconcelos. Robust audio-visual instance discrimination. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12934–12945, 2021. 4
- [39] 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. 6
- [40] Utkarsh Ojha, Yuheng Li, and Yong Jae Lee. Towards universal fake image detectors that generalize across generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 24480–24489, 2023. 3
- [41] KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Namboodiri, and CV Jawahar. A lip sync expert is all you need for speech to lip generation in the wild. In *Proceedings of the 28th ACM international conference on multimedia*, pages 484–492, 2020. 1, 6
- [42] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 1, 2, 4
- [43] 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. 2, 6, 7
- [44] Andrew Rouditchenko, Angie Boggust, David Harwath, Brian Chen, Dhiraj Joshi, Samuel Thomas, Kartik Audhkhasi, Hilde Kuehne, Rameswar Panda, Rogerio Feris, et al. Avlnet: Learning audio-visual language representations from instructional videos. In *Annual Conference of the International Speech Communication Association*. International Speech Communication Association, 2021. 4
- [45] Ludan Ruan, Yiyang Ma, Huan Yang, Huiguo He, Bei Liu, Jianlong Fu, Nicholas Jing Yuan, Qin Jin, and Baining Guo. Mm-diffusion: Learning multi-modal diffusion models for joint audio and video generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10219–10228, 2023. 8
- [46] Laurens van der Maaten and Geoffrey E. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9: 2579–2605, 2008. 2
- [47] Jun Wang, Yinglu Liu, Yibo Hu, Hailin Shi, and Tao Mei. Facex-zoo: A pytorch toolbox for face recognition. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 3779–3782, 2021. 4
- [48] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10039–10049, 2021. 1
- [49] 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 (CVPR)*, pages 4129–4138, 2023. 3
- [50] Deressa Wodajo and Solomon Atnafu. Deepfake video detection using convolutional vision transformer. *arXiv preprint arXiv:2102.11126*, 2021. 3, 6

- [51] Lior Wolf, Tal Hassner, and Itay Maoz. Face recognition in unconstrained videos with matched background similarity. In *CVPR 2011*, pages 529–534. IEEE, 2011. [2](#)
- [52] Haojie Wu, Pan Hui, and Pengyuan Zhou. Deepfake in the metaverse: An outlook survey. *arXiv preprint arXiv:2306.07011*, 2023. [1](#)
- [53] 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. [2](#), [3](#), [6](#)
- [54] Xin Yang, Yuezun Li, and Siwei Lyu. Exposing deep fakes using inconsistent head poses. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8261–8265, 2019. [3](#)
- [55] Xiaohui Zhao, Yang Yu, Rongrong Ni, and Yao Zhao. Exploring complementarity of global and local spatiotemporal information for fake face video detection. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2884–2888, 2022. [3](#)
- [56] Yinglin Zheng, Jianmin Bao, Dong Chen, Ming Zeng, and Fang Wen. Exploring temporal coherence for more general video face forgery detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 15044–15054, 2021. [3](#), [6](#), [7](#)
- [57] 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. [3](#), [7](#)
- [58] Yang Zhou, Zhan Xu, Chris Landreth, Evangelos Kalogerakis, Subhansu Maji, and Karan Singh. Visemenet: Audio-driven animator-centric speech animation. *ACM Transactions on Graphics (TOG)*, 37(4):1–10, 2018. [1](#)
- [59] Hao Zhu, Wayne Wu, Wentao Zhu, Liming Jiang, Siwei Tang, Li Zhang, Ziwei Liu, and Chen Change Loy. Celebv-hq: A large-scale video facial attributes dataset. In *European conference on computer vision*, pages 650–667. Springer, 2022. [2](#)
- [60] Wanyi Zhuang, Qi Chu, Zhentao Tan, Qiankun Liu, Haojie Yuan, Changtao Miao, Zixiang Luo, and Nenghai Yu. Uia-vit: Unsupervised inconsistency-aware method based on vision transformer for face forgery detection. In *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part V*, page 391–407, Berlin, Heidelberg, 2022. Springer-Verlag. [3](#)