

Avatar Fingerprinting for Authorized Use of Synthetic Talking-Head Videos

Ekta Prashnani Koki Nagano Shalini De Mello David Luebke Orazio Gallo

NVIDIA

{eprashnani, knagano, shalinig, dluebke, ogallo}@nvidia.com

Abstract

Modern generators render talking-head videos with impressive levels of photorealism, ushering in new user experiences such as videoconferencing under constrained bandwidth budgets. Their safe adoption, however, requires a mechanism to verify if the rendered video is trustworthy. For instance, for videoconferencing we must identify cases in which a synthetic video portrait uses the appearance of an individual without their consent. We term this task avatar fingerprinting. We propose to tackle it by leveraging facial motion signatures unique to each person. Specifically, we learn an embedding in which the motion signatures of one identity are grouped together, and pushed away from those of other identities, regardless of the appearance in the synthetic video. Avatar fingerprinting algorithms will be critical as talking head generators become more ubiquitous, and yet no large scale datasets exist for this new task. Therefore, we contribute a large dataset of people delivering scripted and improvised short monologues, accompanied by synthetic videos in which we render videos of one person using the facial appearance of another¹.

1. Introduction

Recent face portrait generators can synthesize real-time talking-head videos hardly distinguishable from real ones. Despite the risks for visual disinformation it poses, the legitimate use of synthetic avatars will become ubiquitous, bringing benefits to a myriad of applications ranging from personalized avatars to AR filters for selfie videos.

In the context of video conferencing, for instance, instead of sending the entire video of a person talking, we can synthesize a video on the receiver’s end using a single frame to specify the target identity, together with some compact information about the motion of the driving video, thus saving valuable bandwidth [43]. To enable safe use in such cases, the relevant question is no longer whether a video is “real” or not, but whether it is “trustworthy” or not. That is,

¹Project page: <https://research.nvidia.com/labs/nxp/avatar-fingerprinting/>




Figure 1: Talking head generators produce realistic videos of a target identity from driving videos of different identities. Our method extracts identity-agnostic temporal facial features and learns an embedding in which videos driven by one identity fall close to each other and far from those driven by other identities, regardless of the appearance.

we want to determine if the person driving the video (ID₁ or ID₂ in Figure 1) is authorized to control the likeness, or the appearance, of the synthetic video portrait (target identity in Figure 1). We call this novel task *avatar fingerprinting*.

To perform avatar fingerprinting, we leverage a simple but fundamental observation: each person shows unique facial motion idiosyncrasies when talking and emoting. For instance, someone may raise one of her eyebrows more than the other, or shake their head more often while talking. These “dynamic identity signatures” [38] have been shown to carry enough information for humans to recognize other individuals, *even when the physical appearance of their face is altered* [38, 22, 30]. This makes them attractive for our task, as they are derived solely from the driving identity of an avatar, regardless of the appearance of the generated video.

Fortunately, the level of sophistication of modern talking-head generators allows them to capture even subtle facial motions—and in the future they will improve even further. To leverage this observation, then, we extract facial landmarks and their temporal dynamics from the video. We then use a novel contrastive loss that allows us to uniquely




Figure 2: We introduce a large dataset of real and synthetic talking-head videos. We capture 46 subjects talking in both scripted and free-form settings. To encourage a natural performance in a realistic setting, we record the subjects while videoconferencing with each other (left). We also synthesize talking-head videos by using driving identities from both our own data and other datasets (see Section 4.2). In the synthetic videos on the right, each row corresponds to a driving identity ($ID_i \rightarrow (\cdot)$) and each column corresponds to a different target identity ($(\cdot) \rightarrow ID_j$). The videos in which driving and target identity match are **self-reenactments** videos, the rest are **cross-reenactments**.

fingerprint them using a *dynamic identity embedding* space of the talking head (see Figure 1). In this space, the embeddings corresponding to one driving identity across multiple videos and target identities stay close to each other, and far from those of other driving identities.

Avatar fingerprinting is a new task, and no existing datasets completely serve its training and validation requirements. For a comprehensive evaluation of our algorithm, we desire videos of multiple subjects delivering both scripted and natural monologues, captured under typical conditions, such as varying video quality and gaze direction. We also need synthetic talking-head videos in which driving and target identities are different (cross-reenactment), and the case in which they match (self-reenactment). Unfortunately, existing datasets of real videos provide either scripted [33] or natural [46] monologues, and synthetic video datasets focus either on self-reenactment [18] or on cross-reenactment [31, 39], but not both. Therefore, we collected a large dataset of 46 subjects delivering both scripted and free-form monologues, complete with synthetic videos of self- and cross-reenactment [43]. We think that authenticating the identity of a synthetic video regardless of the target identity will be increasingly critical and hope to aid community progress in this direction by releasing our dataset. Therefore, we took particular care in designing the capture protocol, including its legal and privacy-related aspects. In summary:

- We introduce the task of avatar fingerprinting, which focuses on verifying the driving identity of synthetic talking-head videos, rather than detecting them as synthesized.
- We release the first large dataset of subjects delivering scripted and natural monologues, complete of self- and cross-reenactment synthetic videos.

- We propose a baseline solution to this novel task in the context of video conferencing by extracting person-specific motion signatures.

2. Related Work

Our avatar fingerprinting task addresses the authentication of synthetic media, but we do not focus on active forensics techniques such as watermarking. While numerous efforts exist in detecting deepfakes, to our knowledge, no previous work has addressed authorized use of synthetic talking-head videos as we do. While we do not focus on the detection of deepfakes, our method is closely related to a family of deepfake detection techniques that operate based on person-specific identity features, such as the behavioral features of a face. While we focus on verifying talking head videos, we do not assume anything specific about the talking head generators and the same algorithm can be applied to other types of generators in the future.

Learning-based Attribution of Synthetic Media.

Learning-based approaches have been used to learn patterns and common characteristics in images or videos to identify the origin and authenticity of the media, or to determine if it has been manipulated or altered in some way. Previous work [6, 26] used a pre-trained GAN generator to attribute a synthesized image to its generator by projecting the image to the generator’s latent space. In this case, a perfectly inverted image indicates the image is generated by the generator, while a real image is less invertible. Based on this, GAN-Scanner [47] used a variant of StyleGAN2-based inversion to detect images generated by never-before-seen GANs. A learning-based approach was also used to learn camera fingerprints associated with videos recorded by physical cameras to determine whether

a media is manipulated [15]. Another line of work to verify the authenticity of media is by embedding imperceptible watermarks into images and videos [19, 7, 42, 34]. A previous work [48] showed that they can embed invisible watermarks into training data and a GAN model trained with it can successfully transfer watermarks to the generated images, allowing attribution of the generator. More recent work [49] showed a framework to embed watermarks in the form of a conditional GAN model for scalable fingerprinting of generated images. While these are viable solutions, our method focuses on a *passive* forensics technique that operates given a probe video and a set of verified user videos.

Deepfake Detection Based on Identity-Specific Features.

A family of deepfake detectors use identity-specific features to verify whether the person depicted in a video is who it purports to be, by posing the detection problem as an identity-recognition problem. The advantage of this approach is that the forensics method only needs to be trained with pristine videos and does not need to be trained against specific generators, enabling generalization to unseen generators. Agarwal *et al.* [4] is the first to exploit person-specific patterns in facial expressions and poses as the distinctive identity signature to detect fake videos. Follow up works explored other types of person-specific soft biometric features, such as, vocal mannerisms [10], phoneme-viseme consistencies [3], word-facial expression consistencies [5], and dynamics of ears [2]. While many of these works need person-specific training, previous works [32, 1, 14, 13] extended this idea to train a CNN-based detector using a large-scale in-the-wild video data [12] and variants of contrastive learning [40, 44, 20, 28]. Agarwal *et al.* [1] combined static facial appearance using a facial recognition model and dynamic facial behaviors using a CNN, and showed that this approach is effective for detecting face-swap deepfakes. ID-Reveal [14] used facial shapes and motions encoded in a low-dimensional space of a 3D morphable model [8] and demonstrated the ability to handle both face-swapping and talking-head deepfakes. Beyond these dynamic facial identity features, previous works explored temporal inconsistencies of face identities within a video [32] and identity consistencies of inner and outer face regions [16]. While these detection approaches aim for differentiating real and synthetic videos, our avatar fingerprinting task requires determining a driver *among* the synthetic videos.

Dynamic Facial Identity Signatures. In ongoing research on how humans recognize and attribute faces, cognitive scientists have studied the impact of “dynamic facial identity signatures” (*i.e.*, characteristic or identity-specific movements of a face) for identity recognition, especially

when limited appearance-specific “static identity” cues are available [38]. In one experiment, scientists projected facial animations generated by human actors onto a computer-generated average head [22]. Subjects were able to learn to discriminate individuals based solely on facial motion. In another, subjects correctly attributed animations of synthetic faces and to their morphed versions [30]. These studies have focused on the learnability of dynamic signatures by human observers and observed a positive trend when dynamic behavior was the most or the only reliable cue to identity. Most recent work performed human user studies in the context of deepfake videos and suggested that behavior signatures could be used to distinguish synthesized videos regardless of their visual appearance [37]. We wish to seek such “dynamic facial identity signatures” in our efforts to fingerprint a synthetic talking head. This will allow us to abstract away the appearance of an avatar, and therefore, attribute an avatar to the driving identity instead of the target identity shown in the avatar.

Talking Head Generators. Given a target portrait and a driving video of a user, modern talking head generators [45, 24, 41, 51, 50, 27, 17, 43, 25] reenact the target portrait by using motions (facial expressions and head poses) from the driving video. Instead of using a single target frame, talking head generators may use an image set to capture better person-specific appearance and unseen background [35]. To transfer the motions from the driver video, most of these generators rely on additional deep neural networks to decompose motions from the appearance of the driving video in a form of appearance-agnostic sparse keypoints and dense motion fields.

3. Terminology

We seek to verify the trustworthiness of a synthesized talking-head video, termed *target video*. We assume that an avatar-generation tool (*e.g.*, [43]) created it by animating an image (*target image*) using the expressions and head poses obtained from another video, which we term the *driving video*. We call *driving identity* the identity of the person in the driving video, and *target identity* the identity of the person in the target image. When driving and target identities match, the target video is a *self-reenactment*, while the case of a driving identity used to animate a different target identity is *cross-reenactment*. This allows us to more formally state our goal: we want to verify that a target video is self-reenactment. With this terminology, we introduce our dataset, which includes real videos as well as self- and cross-reenactment videos.

4. Dataset

Recall that avatar fingerprinting is not about detection of synthetic media. Rather, we already know a video to

Dataset	Subjects	Free or Scripted	Performance	(S)elf- or (C)ross-reenact.?
RAVDESS [33]	24	Scripted	Prescribed	Real only
CREMA-D [11]	91	Scripted	Prescribed	Real only
VFHQ [18]	36	Scripted	Prescribed	Real + (S) of [33]
CDFv2 [31]	59	Free	Natural	Real + (C)
Ours	46	Both	Natural	Real + (S) + (C)

Table 1: Comparison with existing datasets. “Performance” refers to whether the subjects were asked to use a prescribed emotion (*e.g.*, anger). Our dataset is the first to have both scripted and free-form monologues, and self- and cross-reenactments to accompany the real videos.

be synthesized and seek to verify the driving identity, to enable its safe use. Therefore, we need a dataset comprising of both self- and cross-reenacted videos. Alas, existing datasets only offer real, self-reenacted, or cross-reenacted videos, but not all at the same time, see Table 1. We introduce a large dataset of real and synthetic talking-head videos that fills this void. Figure 2 shows an overview of how the videos are captured and synthesized.

4.1. Real Data Capture

Capturing videos of monologues delivered by different subjects for the purpose of identity verification introduces two conflicting goals. On the one hand a controlled evaluation of the trained models requires predictability of what is spoken to prevent identification algorithms from latching onto the spoken content itself. On the other, we want the subjects to act as they would in a casual conversation, rather than with a prescribed emotion, to capture their uniquely identifying mannerisms. We address this tradeoff by recording the subjects while videoconferencing with each other, which creates the impression of being in a natural conversation. This differs from existing datasets, in which the subjects are recorded while looking at the camera. We also design two distinct recording strategies: a free-form stage where the subjects are given only general guidance on the topics, and a more controlled scripted stage, see Figure 2(a). This paradigm fits well with our main application, which is avatar fingerprinting for video conferencing. We captured our videos with minimal instructions on how to setup the video conference, allowing for the variability one can expect in a natural setting. For example, the backgrounds present various degrees of clutter, and the available bandwidth often affects the video compression. This was by design: we want the dataset to be as challenging as real-life scenarios. In total we record 46 subjects of diverse genders, age groups, and ethnicity.

Stage I: Free-Form Monologues. In this first stage, the two subjects on the call alternate between asking and answering seven pre-defined questions. The questions are designed to avoid sensitive or potentially inflammatory topics.

This is critical because we later use sentences spoken by one individual to animate the video of a second individual, quite literally putting words in their mouths. Questions include topics such as one’s favorite family holiday, or their least favorite house chore. (The complete list of questions is in the Supplementary.) To further create a natural interaction, the subject listening is encouraged to actively engage with the one speaking (*e.g.*, by nodding or smiling), while remaining silent.

Stage II: Scripted Monologues. For this stage, we prepared thirty short utterances consisting of two or three sentences each. We chose this length to allow for memorization, while still providing enough content to trigger facial expressions. However, to avoid inducing unnatural expressions, we did not prescribe specific emotions for each utterance. For instance, we did not ask to express anger for a sentence, but we did choose sentences that may naturally evoke it, and used punctuation to encourage it, *e.g.* “Will you please answer the darn phone? The constant ringing is driving me insane!” We instructed the subjects to split their screens to show both this list and the call video, see Stage II in Figure 2(a). More details, including the full list of utterances can be found in the Supplementary.

Privacy Considerations. Face videos are sensitive data, since a person’s face is a key identifier. We took on this task with care to ensure good data governance. Our proposal for the data capture protocol was approved by an Institutional Review Board (IRB). Our goal was to provide exhaustive and transparent information to participants about our data capture procedure, future plans with the dataset (including our intent to create synthetic data samples), and conditions under which future research would be conducted—by us and interested third parties. The participants were also asked to confirm whether their data can be used for research beyond identity validation, and whether it could be shown in public disclosures. Each file in our dataset is annotated with their responses.

4.2. Synthetic Talking-Head Videos

In addition to the videos described in Section 4.1, we also need synthetic videos to train and evaluate our avatar fingerprinting algorithm. Existing datasets focus either on self-reenactment [18] or on cross-reenactment [31, 39], but not both. Our dataset fills this void and even uses identities from one dataset to synthesize target identities from different datasets.

Specifically, we pool the 91 identities from the original videos of CREMA-D [11], the 24 identities from those of RAVDESS [33], and the 46 from our own dataset, for a total of 161 unique identities $\mathcal{I}$. Recall that we have several real videos for each identity $ID_i \in \mathcal{I}$. To avoid a combinatorial




Figure 3: We extract landmarks from the frames of a talking-head clip, compute their normalized pairwise distances, and concatenate the frame-wise features. We then learn an identity embedding using a loss that pulls closer features of videos driven by the same identity and pushes away those driven by others. $ID_i \rightarrow ID_j$ indicates a video of identity j driven by identity i .

explosion of synthetic videos, for all pairs of identities ID_i and ID_j , we use ID_j as the target identity and we randomly select 8 of the videos of ID_i to generate 8 cross-reenactment videos, $\{ID_i^k \rightarrow ID_j\}_{k=\{1,\dots,8\}}$ (all 8 share the same target image). We also generate a self-reenactment video for each original video. We use Face-vid2vid for synthesizing the videos [43]. We chose this portrait generator because it is the state-of-the-art and it preserves the identity-specific facial motion dynamics, as needed to perform avatar fingerprinting. Nevertheless, the reconstruction is not perfect; for instance, in the third row of Figure 2(b) the person’s eyes are squinted in the driving video (ID_3), but completely closed in all reenactments—including in self-reenactment. In total we generate $\sim 218,000$ cross-reenactment videos. More details are in the Supplementary.

5. Method

Overview. We seek to verify the driving identity of a synthetic video, independently of the target identity. We leverage the finding from cognitive science research that each person emotes in unique ways when communicating, and that this signal is sufficient for recognition, even when the actual appearance is artificially corrupted [38, 22, 30]. We note that these dynamic features are not a bi-product of the generator, and come from the driving identity itself: they may be the way a person smiles, or the way she frowns. Notably, they are distinct from the temporal artifacts introduced by the generator, and that existing algorithms use to detect whether a video is synthetic or real [21].

An overview of our algorithm for avatar fingerprinting is shown in Figure 3. To capture these expressions, we extract the relative position of facial landmarks over time from the input video, as shown in Figure 3(a) (Section 5.1). We then project these temporal signatures onto a *dynamic identity*

embedding in which features belonging to the same driving identity are close to each other regardless of the target identity used to generate the video, *i.e.*, independently of appearance (Figure 1). To learn this embedding we train a neural network with a novel contrastive loss that pulls together all embedding vectors of synthetic videos driven by an individual, while pushing away the embedding vectors of videos driven by all others individuals (Section 5.2). More details on the implementation are in Section 5.3.

5.1. Dynamic Facial Identity Features

Our first step is to extract temporal features that summarize short segments of the video we wish to fingerprint. For each frame we want to extract features that must:

1. have minimal dependency on the appearance of the face in the video (that is, the target identity),
2. reflect the dynamics of the expressions, and
3. can capture subtle expressions.

One choice could be per-frame 3DMM features [9]: a strategy also used by Cozzolino *et al.* to detect synthetic videos [14]. However, we empirically observe that 3DMM features are not sufficiently expressive, and do not satisfy requirement 3 (see ablation experiments in Section 6). We observe a similar behavior for action units [4]. Facial landmarks [23] address this issue, but are sensitive to the shape of the face in the video, and thus to the target identity.

To leverage the expressiveness of facial landmarks while abstracting from the underlying facial shape, we propose to compute the pairwise normalized Euclidean distance between each of the landmarks of a frame. We concatenate these distances into a single vector for the frame, d_f . A small subset of the landmark features and their distances are shown in Figure 3(a).

We then break the input video into *clips*, which are sequences consisting of F frames and offset by one frame (*e.g.*, $[1, F]$, $[2, F+1]$, *etc.*), and concatenate the vectors from all the frames in each clip. Using the change in the relative position of the landmarks over a short period of time (the length of a clip) allows us to capture temporal dynamics with minimal dependence on the absolute position of each landmark, *i.e.* independently of the shape of the face.

We show empirically that our features are a good representation for our task, and can even improve the results of baseline state-of-the-art methods (Section 6).

5.2. Dynamic Identity Embedding Contrastive Loss

While the features described in Section 5.1 extract low-level motion dynamics, they cannot be used directly to disambiguate two target videos based on the driving identity. We tackle this problem by learning a dynamic identity embedding, a space where videos driven by one subject map to points that are close to each other and far from the videos driven by anybody else.

Specifically, we use a temporal CNN to extract an embedding vector from a clip, which, as described before, is a short segment of an input video. To train the network we use a dataset of synthetic videos driven by different identities. We denote as $\mathcal{C}_{\text{ID}_1 \rightarrow \text{ID}_2}^k(t)$ the embedding produced by the network for the clip starting at time t in the k -th video, of a target identity ID_2 driven by identity ID_1 . As stated above, we have two main objectives, which we capture with the following two terms in our proposed loss function.

We Want to Pull Together All the Videos Driven by ID_1 . To achieve this, we define the following term:

$$N_{j,\text{ID}_1,\text{ID}_2}(t) = \sum_{\text{ID}_l,k} \max_n s(\mathcal{C}_{\text{ID}_1 \rightarrow \text{ID}_2}^j(t), \mathcal{C}_{\text{ID}_1 \rightarrow \text{ID}_l}^k(n)), \quad (1)$$

where $s(\cdot, \cdot) = e^{-\|\cdot\|^2}$ is a similarity metric. Intuitively, Equation 1 takes two videos, j and k , both driven by ID_1 . Given a clip starting at time t in the first video, it looks for the most similar clip in the second video. Since the driving identity is the same for both videos, Equation 1 encourages an embedding where clips that capture a similar expression are closer to each other. Equation 1 is high even if only one clip from video k has a similar temporal signature to $\mathcal{C}_{\text{ID}_1 \rightarrow \text{ID}_2}^j(t)$. That is because even just one occurrence of the same expression is evidence that the driving identity may be the same. Of course other driving identities may use similar expressions and we address that with the loss term described below. Additionally, we note that k spans the set of *all* videos driven by ID_1 , and ID_l spans *all* identities, including $\text{ID}_l = \text{ID}_1$ and $\text{ID}_l = \text{ID}_2$.

We Want to Push Away Videos not Driven by ID_1 . We define the following term:

$$Q_{j,\text{ID}_1,\text{ID}_2}(t) = \sum_{\text{ID}_l \neq \text{ID}_1,k} \max_n s(\mathcal{C}_{\text{ID}_1 \rightarrow \text{ID}_2}^j(t), \mathcal{C}_{\text{ID}_l \rightarrow \text{ID}_2}^k(n)), \quad (2)$$

where, similarly to Equation 1, we take a clip from video j , and look for the most similar clip in video k . However, this time the two videos share the same target identity, but are driven by different identities: we want all the videos driven by identities different from ID_1 to be pushed away from those driven by ID_1 , including videos where ID_1 is the target identity. Note that ID_2 spans *all* identities, including $\text{ID}_2 = \text{ID}_1$ and $\text{ID}_2 = \text{ID}_l$.

Combining Equations 1 and 2, we write the probability that the embedding vector $\mathcal{C}_{\text{ID}_1 \rightarrow \text{ID}_2}^j(t)$ lies close to the embedding vectors for *all* video clips driven by ID_1 and far from *all* the videos driven by others as

$$p_{j,\text{ID}_1,\text{ID}_2}(t) = \frac{N_{j,\text{ID}_1,\text{ID}_2}(t)}{N_{j,\text{ID}_1,\text{ID}_2}(t) + Q_{j,\text{ID}_1,\text{ID}_2}(t)}, \quad (3)$$

and the complete loss term as

$$\mathcal{L} = \sum_{j,\text{ID}_1,\text{ID}_2,t} -\log(p_{j,\text{ID}_1,\text{ID}_2}(t)). \quad (4)$$

5.3. Implementation

Parameter Choices. To extract the per-frame dynamic facial identity features $\mathbf{d}_f$, we first estimate 126 facial landmarks for each frame using a pre-existing implementation [23]. We compute the Euclidean distances between all possible pairs of landmarks, to obtain a 7875-dimensional vector for each frame. For all of the 161 identities in the synthetic component of our dataset, we scan through original videos to isolate a frame for each subject that shows a neutral head pose and expression (frontal head pose, no expression). The square root of the area occupied by the bounding box of their face in this neutral frame is used to normalize the 7875 values of the pairwise distance vector, resulting in $\mathbf{d}_f$. The clip duration F is set to 51. A shorter clip, while allowing computational efficiency, might not capture meaningful expressions. In contrast, longer clips might capture multiple expressions making it tough to match to other clips. After experimenting with different values for F (see supplementary), we chose $F = 51$ as it allowed us to learn a meaningful dynamic identity signature.

Training Details. Temporal convolutional networks that operate on intermediate representations of frames have shown remarkable success in modeling facial behavior and its anomalies [21, 36, 14]. We use one such architecture, the temporal ID network [14], after adapting the input layer to match our input feature dimensions and, when needed, input clip duration. This input tensor is obtained by concatenating $\mathbf{d}_f$ across F frames of the clip, to obtain a tensor of size $7875 \times F$. The neural network outputs a 128-dimensional embedding vector for each clip, which is trained to cluster based on the driving identity (Section 5.2). In each batch, we include 8 unique identities. For each identity ID_i , the pull term (Equation 1) comprises of 16 clips: 8 are self-reenactments, randomly sampled from the full set, and the remaining are cross-reenactments with ID_i as the *driving* identity. These cross-reenactments can potentially show the same words being spoken by different target identities. This is crucial: it allows the neural network to learn to pull together videos based purely on the facial motion, regardless of the appearance of the video. The push term (Equation 2) for ID_i is composed of clips with the remaining 7 identities in the batch serving as driving identities (8 clips per driving identity). Therefore, for each identity, 72 clips are included in a batch. The training is performed for 100,000 iterations, with Adam optimizer [29] and a learning rate of $1e^{-4}$. We will discuss additional details in the supplementary.




Figure 4: ROC curves for our method and two baselines.

Training, Validation, and Test Datasets. We split the training, validation, and test datasets based on identities. Out of the total 161 identities (pooling together the identities from our dataset, RAVDESS, and CREMA-D— see Section 4.2), we reserve 35 for testing, 14 for validation and 112 for training. We ensure that there is no cross-set cross-reenactments: that is, identities in the training set only drive other training-set identities (and similarly for validation and test set). This allows us to evaluate the generalizability of our method to novel facial mannerisms that are not seen during training, as well as novel target identities.

6. Evaluation

We begin by evaluating qualitatively our method’s ability to extract embedding vectors based on the driving identity. Figure 5 shows a set of self- and cross-reenacted clips (please view the animation in a media-enabled viewer). For each row, we take one identity as reference and we compute the embedding vectors of clips that use it both as driving and target identity. We then compute the average Euclidean distance of the resulting embedding vectors against those of *all other clips driven by the same reference identity*. We note that the average distance d is lower when the driving identity matches the reference identity (first two columns). We also note that the distance between the clips in the first two columns is similar: this confirms that it is a function of the motion, rather than the appearance. When the driving identity changes, the average distance increases, even if the target identity matches the reference identity, which is precisely our goal. More results are in the Supplementary.

To evaluate our approach more formally, we use 35 unique identities not used as driving or target identities in the training set. One at a time, we treat each identity ID_i as target and synthesize cross-reenactments using all the remaining identities as drivers. This is the set of “unauthorized” synthetic videos for ID_i . The self-reenacted samples for ID_i form the “authorized” set. Note that there are several self-reenacted videos of ID_i , one per original video of ID_i .

For each target identity ID_i , we extract the dynamic identity embedding vector of all the clips in the pool of its self-




Figure 5: **Animated figure. Open in a media-enabled viewer like Adobe Reader and click on the inset.** Our embeddings capture the dynamics of an expression, rather than the appearance of the face. For each row, we pick a reference identity. The **green box** indicates reenactments driven by the reference identity, the **red** and **blue** are cross-reenactments of the reference identity. We compute the average distance of each clip shown here against all other clips driven by the reference identity. The average distance to the other clips of the reference identity is consistent for a given motion, and lower (better) when the reference identity is driving as compared to the cross-reenactments that use the reference identity as target.

and cross-reenacted videos, and compute their Euclidean distances. That is, for clip k we compute

$$d(C_{ID_i \rightarrow ID_i}^k, C_{ID_i \rightarrow ID_i}^l), \quad \forall l \neq k, \text{ and} \quad (5)$$

$$d(C_{ID_i \rightarrow ID_i}^k, C_{ID_j \rightarrow ID_i}^l), \quad \forall l \neq k, \quad \forall i \neq j. \quad (6)$$

We threshold these distances for each target identity to get an ROC curve, and average across the target identities to get the overall area under the curve (AUC). We note that this AUC measures one model’s ability to classify a synthetic video as self-reenactment or as cross-reenacted.

6.1. Comparisons with Existing Methods

Avatar fingerprinting is a novel task, and no existing methods directly address it. The closest related works aim at detecting real versus synthetic media. As discussed in Section 2, some of these detectors learn identity-specific features such as facial expressions and head poses [4], or facial shapes and motion [14] and can serve as baselines for the task of avatar fingerprinting with some adaptation. The work by Agarwal *et al.* trains a model to detect synthetic videos of a *specific* identity [4]. To adapt it to our task, we train 35 different models, one for each identity

in the evaluation, by splitting the corresponding original videos in two subsets. We then test each model on the self- and cross-reenactment videos of the corresponding identity. ID-Reveal, trained on a large-scale dataset, learns an embedding space where real videos of a specific identity are grouped together [14]. Since the method shows good generalization to new identities, for the task of synthetic media detection, we directly use the pre-trained model on our data to detect, once again, self- versus cross-reenactment. Figure 4 shows the ROC curve for our method compared to these baselines. Our method (AUC=0.810) outperforms by a wide margin both ID-Reveal (AUC=0.720), and the method by Agarwal *et al.* (AUC=0.726). We also note that, unlike ID-Reveal and our method, Agarwal *et al.* uses a different model for each identity.

6.2. Ablation Study

Our method outperforms existing baselines by introducing two novel components: the dynamic facial identity features, which capture in a compact and expressive way the face dynamics, and the loss function, which defines the shape of the identity embedding. Here we study the contribution of each. We evaluate the contribution of our dynamic facial identity features by swapping them with 3DMM features [8], a popular choice to capture facial dynamics. Since we use ID-Reveal as our temporal CNN backbone, for this ablation we use the loss function proposed in their original paper [14]. We re-train the same network using our features and observe a jump from 0.677 to 0.727 in terms of AUC. Upon inspection we notice that the 3DMM features tend to over-smooth the facial motion, and are unable to capture subtle dynamics that prove critical to avatar fingerprinting, and which our features capture (more visual comparison are in the Supplementary). We also evaluate the contribution of our identity embedding loss and observe a further improvement. Table 2 summarizes this ablation study. ID-Reveal’s loss distinguishes between real and synthetic videos. Therefore, it can be satisfied more easily because of the artifacts that only affect synthetic videos. Our loss only compares synthetic videos and, therefore, must latch on the subtle differences in self- and cross-reenactment videos: the underlying dynamics of the motion.

6.3. Limitations

Our algorithm is less discriminative of subjects that are less emotive and more neutral most of the times. In the future, relying on more granular dynamic signatures that can extract micro-expressions can help alleviate this. Since our avatar fingerprinting algorithm relies on facial landmarks, their robustness affects its accuracy as well. The performance of our method degrades when expressions that are critical to verifying the driving identity are not captured by the synthetic portrait generator. For instance, if a synthetic

Input Features	Loss	AUC
3DMM	ID-Reveal rec. loss	0.677
LM dist	ID-Reveal rec. loss	0.727
LM dist	Our loss	0.810

Table 2: Ablation study showing the importance of our input features and loss function.

portrait generator does not animate eyebrows in the rendered video, we cannot discriminate between identities that differ only in the way they move their eyebrows. Lastly, our dataset currently features only one type of interaction: one-on-one conversations. Expanding it to include other forms of conversational interactions would be beneficial.

7. Societal Impact

We acknowledge the societal importance of introducing guardrails when it comes to the use of talking-head generation technology. We present this work as a step towards trustworthy use of such technologies. Nevertheless, our work could be misconstrued as having solved the problem and inadvertently accelerate the unhindered adoption of talking head technology. We do not advocate for this. Instead we emphasize that this is only the first work on this topic and underscore the importance of further research in this area. Since our dataset contains human subjects’ facial data, we have taken many steps to ensure proper use and governance with steps including: obtaining IRB approval, informed subject consent prior to data capture, removing subject identity information, pre-specifying the subject matter that can be discussed in the videos, allowing subjects the freedom to revoke our access to their provided data at any point in future (and stipulating that interested third parties maintain current contact information with us so we can convey these changes to them).

8. Conclusions

Highly photo-real portrait talking-head generators are becoming increasingly beneficial to applications such as video conferencing. This trend raises the important new research question of how best to also ensure their safe use in such scenarios. To this end, we investigate the new problem of avatar fingerprinting to authenticate legitimate talking-heads created by authorized users. We leverage the fact that driving individuals have uniquely identifying dynamic motion signatures, which are also preserved in the videos that they drive. Since none exists, we contribute a new large dataset carefully designed to further research on avatar fingerprinting. While our seminal work shows promising results and establishes a baseline in the field, we hope that it lays the foundation for much further research on this deeply impactful topic.

Acknowledgements. We would like to thank the participants for contributing their facial audio-visual recordings for our dataset, and Desiree Luong, Woody Luong, and Josh Holland for their help with Figure 2. We thank David Taubenheim for the voiceover in the demo video for our work, and Abhishek Badki for his help with the training infrastructure. We acknowledge Joohwan Kim, Rachel Brown, Anjul Patney, Ruth Rosenholtz, Ben Boudaoud, Josef Spjut, Saori Kaji, Nikki Pope, Daniel Rohrer, Rochelle Pereira, Rajat Vikram Singh, Candice Gerstner, Alex Qi, and Kai Pong for discussions and feedback on the security aspects, data capture and generation steps, informed consent form, photo release form, and agreements for data governance and third-party data sharing. Koki Nagano, Ekta Prashnani, and David Luebke were partially supported by DARPA’s Semantic Forensics (SemaFor) contract (HR0011-20-3-0005). This research was funded, in part, by DARPA’s Semantic Forensics (SemaFor) contract HR0011-20-3-0005. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. Distribution Statement “A” (Approved for Public Release, Distribution Unlimited).

References

- [1] Shruti Agarwal, Tarek El-Gaaly, Hany Farid, and Ser-Nam Lim. Detecting deep-fake videos from appearance and behavior. *2020 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–6, 2020. 3
- [2] Shruti Agarwal and Hany Farid. Detecting deep-fake videos from aural and oral dynamics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2021. 3
- [3] 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 (CVPR) Workshops*, 2020. 3
- [4] Shruti Agarwal, Hany Farid, Yuming Gu, Mingming He, Koki Nagano, and Hao Li. Protecting world leaders against deep fakes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2019. 3, 5, 7
- [5] 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 *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2023. 3
- [6] Michael Albright and Scott McCloskey. Source generator attribution via inversion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2019. 2
- [7] Shumeet Baluja. Hiding images in plain sight: Deep steganography. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 3
- [8] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of SIGGRAPH*, 1999. 3, 8
- [9] Volker Blanz and Thomas Vetter. Face recognition based on fitting a 3D morphable model. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2003. 5
- [10] Matyáš Boháček and Hany Farid. Protecting world leaders against deep fakes using facial, gestural, and vocal mannerisms. *Proceedings of the national academy of Sciences*, 2022. 3
- [11] Houwei Cao, David G Cooper, Michael K Keutmann, Ruben C Gur, Ani Nenkova, and Ragini Verma. CREMA-D: Crowd-sourced emotional multimodal actors dataset. *IEEE Transactions on Affective Computing*, 2014. 4, 12
- [12] J. S. Chung, A. Nagrani, and A. Zisserman. Voxceleb2: Deep speaker recognition. In *INTERSPEECH*, 2018. 3
- [13] Davide Cozzolino, Matthias Nießner, and Luisa Verdoliva. Audio-visual person-of-interest deepfake detection, 2022. 3
- [14] Davide Cozzolino, Andreas Rössler, Justus Thies, Matthias Nießner, and Luisa Verdoliva. ID-Reveal: Identity-aware DeepFake video detection. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 3, 5, 6, 7, 8, 12
- [15] Davide Cozzolino Giovanni Poggi Luisa Verdoliva. Extracting camera-based fingerprints for video forensics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2019. 3
- [16] 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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [17] Nikita Drobyshev, Jenya Chelishev, Taras Khakhulin, Aleksei Ivakhnenko, Victor Lempitsky, and Egor Zakharov. MegaPortraits: One-shot megapixel neural head avatars. 2022. 3
- [18] Gereon Fox, Wentao Liu, Hyeonwoo Kim, Hans-Peter Seidel, Mohamed Elgharib, and Christian Theobalt. Video-ForensicsHQ: Detecting high-quality manipulated face videos. In *IEEE International Conference on Multimedia and Expo*, 2021. 2, 4
- [19] Jessica Fridrich. *Steganography in Digital Media: Principles, Algorithms, and Applications*. Cambridge University Press, 2009. 3
- [20] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality reduction by learning an invariant mapping. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2006. 3
- [21] Alexandros Haliassos, Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Lips don’t lie: A generalisable and robust approach to face forgery detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 5, 6
- [22] Harold Hill and Alan Johnston. Categorizing sex and identity from the biological motion of faces. *Current Biology*, 2001. 1, 3, 5

- [23] Sina Honari, Pavlo Molchanov, Stephen Tyree, Pascal Vincent, Christopher Pal, and Jan Kautz. Improving landmark localization with semi-supervised learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 5, 6
- [24] Fa-Ting Hong, Longhao Zhang, Li Shen, and Dan Xu. Depth-aware generative adversarial network for talking head video generation. 2022. 3
- [25] Hui Zhang Jian Zhao. Thin-plate spline motion model for image animation. 2022. 3
- [26] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [27] Taras Khakhulin, Vanessa Sklyarova, Victor Lempitsky, and Egor Zakharov. Realistic one-shot mesh-based head avatars. In *European Conference on Computer Vision (ECCV)*, 2022. 3
- [28] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 3
- [29] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv*, 2014. 6
- [30] Barbara Knappmeyer, IM Thornton, and HH Bülthoff. Facial motion can determine facial identity. *Journal of Vision*, 2001. 1, 3, 5
- [31] Yuezun Li, Pu Sun, Honggang Qi, and Siwei Lyu. Celeb-DF: A large-scale challenging dataset for DeepFake forensics. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 4
- [32] Baoping Liu, Bo Liu, Ming Ding, Tianqing Zhu, and Xin Yu. Ti2net: Temporal identity inconsistency network for deepfake detection. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2023. 3
- [33] Steven R Livingstone and Frank A Russo. The ryerson audio-visual database of emotional speech and song (RAVDESS): A dynamic, multimodal set of facial and vocal expressions in north american english. *PloS one*, 2018. 2, 4, 12
- [34] Xiyang Luo, Ruohan Zhan, Huiwen Chang, Feng Yang, and Peyman Milanfar. Distortion agnostic deep watermarking. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [35] Arun Mallya, Ting-Chun Wang, and Ming-Yu Liu. Implicit Warping for Animation with Image Sets. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 3
- [36] Brais Martinez, Pingchuan Ma, Stavros Petridis, and Maja Pantic. Lipreading using temporal convolutional networks. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020. 6
- [37] Qiaomu Miao, Sinhwa Kang, Stacy Marsella, Steve DiPaola, Chao Wang, and Ari Shapiro. Study of detecting behavioral signatures within deepfake videos, 2022. 3
- [38] Alice J O’Toole, Dana A Roark, and Hervé Abdi. Recognizing moving faces: A psychological and neural synthesis. *Trends in Cognitive Sciences*, 2002. 1, 3, 5
- [39] Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. FaceForensics++: Learning to detect manipulated facial images. In *IEEE International Conference on Computer Vision (ICCV)*, 2019. 2, 4
- [40] Florian Schroff, Dmitry Kalenichenko, and James Philbin. FaceNet: A unified embedding for face recognition and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 3
- [41] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 3
- [42] Matthew Tancik, Ben Mildenhall, and Ren Ng. Stegastamp: Invisible hyperlinks in physical photographs. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [43] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 3, 5, 11
- [44] Xun Wang, Xintong Han, Weilin Huang, Dengke Dong, and Matthew R Scott. Multi-similarity loss with general pair weighting for deep metric learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [45] Yaohui Wang, Di Yang, Francois Bremond, and Antitza Dantcheva. Latent image animator: Learning to animate images via latent space navigation. In *International Conference on Learning Representations (ICLR)*, 2022. 3
- [46] Lior Wolf, Tal Hassner, and Itay Maoz. Face recognition in unconstrained videos with matched background similarity. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011. 2
- [47] Yaser Yacoob. Gan-scanner: A detector for faces of stylegan+, 2021. <https://github.com/yaseryacoob/GAN-Scanner>. 2
- [48] Ning Yu, Vladislav Skripniuk, Sahar Abdelnabi, and Mario Fritz. Artificial fingerprinting for generative models: Rooting deepfake attribution in training data. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 3
- [49] Ning Yu, Vladislav Skripniuk, Dingfan Chen, Larry Davis, and Mario Fritz. Responsible disclosure of generative models using scalable fingerprinting. In *International Conference on Learning Representations (ICLR)*, 2022. 3
- [50] Egor Zakharov, Aleksei Ivakhnenko, Aliaksandra Shysheya, and Victor Lempitsky. Fast bi-layer neural synthesis of one-shot realistic head avatars. In *European Conference on Computer Vision (ECCV)*, 2020. 3
- [51] Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. Few-shot adversarial learning of realistic neural talking head models. In *IEEE International Conference on Computer Vision (ICCV)*, 2019. 3

Avatar Fingerprinting for Authorized Use of Synthetic Talking-Head Videos (Supplementary)

A9. Dataset

We now provide additional details for our proposed dataset, including details about question prompts and sentences spoken in both stages, instructions to the participants, demographics of the dataset, and other relevant statistics.

General Instructions to Subjects. The subjects were asked to join pre-assigned Google Meet video calls using a laptop or a desktop. For the recorded video call, the subjects were also asked to position themselves so that their face was centered and parallel to the screen. However, in some cases with specific video-conferencing setups, this constraint was only approximately satisfied. Additionally, subjects were instructed to avoid hand motion since it can occlude their face, and also excessive body motion that might impair the visibility of their face. Before beginning each monologue, subjects were asked to speak “start topic” in a loud, clear voice, and, similarly, the end of each monologue was marked by the subjects speaking “end topic”. These keywords allowed for quicker time-stamped transcription to isolate relevant portions of the Google Meet recordings. Right after a subject said “start topic”, they were instructed to pause for a few seconds and look directly at the camera with a frontal head pose, while holding a neutral expression. These frames with neutral expressions are crucial for successful generation of synthetic talking-head videos using face-vid2vid [43]. Face-vid2vid works by transferring expression changes from a driving video to the target image. Therefore, it is important that the expression of the target image and that of the first frame of the driving video match. Asking subjects to provide a neutral expression before commencing with their monologues proves to be an effective way to achieve this: these neutral frames serve as good target images, while driving videos that start with these neutral frames allow for effectively animating the target image showing a similar expression. During the second stage of the data capture, where we record scripted monologues, subjects were instructed to memorize and speak the sentences to their recording partner, without referring back to the printed text from which they memorized the sentences. In case the subject forgot a sentence, they were instructed to start from the beginning of the sentence set. The whole recording session with both subjects in a video call typically lasted an hour, which also included miscellaneous interactions in between the monologues. The current dataset release excludes such interactions and only focuses on data captured for the two stages (Free-Form Monologues and Scripted Monologues).

Stage I: Free-Form Monologues. Subjects were asked to alternate between speaker and prompter roles. The prompter’s task was to ask each of the following questions to the speaker, and the speaker was instructed to answer these questions in their natural manner.

1. Describe a day when you had to rush to an appointment.
2. Talk about an important milestone you have missed in the past and your feelings about it.
3. What is your favorite family holiday?
4. How is the weather in your area typically?
5. Is there a household chore you don’t like doing?
6. Tell me about an incident that really surprised you.
7. Tell me about an incident that really scared you.

Stage II: Scripted Monologues. The following sentence sets were memorized and recited by each subject (alternating with their recording partner) in the second stage of the data capture. We did not ask subjects to explicitly demonstrate specific emotions for any sentence set. Rather, we chose to allow subjects to perform these memorized sentences in a manner natural to them.

1. My friend has a very cute dog. But, he can be scary when he barks.
2. Will you please answer the darn phone? The constant ringing is driving me insane!
3. My aunt was in the hospital for a week. Unfortunately, she passed away yesterday and I will need some time to grieve.
4. I hate rushing to get to the airport. The stress is too much for me to handle.
5. A slice of cake is the perfect ending to a meal. Wouldn’t you agree?
6. It is going to be great working with you! I am surprised we didn’t connect sooner!
7. You need to take the trash out right now! Your whole apartment smells like rotten eggs!
8. My internet connection is unreliable today. I hope it gets better before my meeting or I will have to call in!
9. I know the deadline is around the corner, but I just don’t have any updates yet, I’m sorry.
10. Why can’t the banker figure out what’s going on? I should have got my money last night!
11. It’s really nice out today. I might go for a walk if I get off work early and the kids aren’t back from school.
12. There is a famous coffee shop around the corner that also serves snacks. Would you like to go tonight?
13. My dog almost got run over by a car today! Thank

- God he is safe!
14. It is getting very cold outside. I feel like having some hot chocolate. Would you like some?
 15. I have been exercising so much lately. But I am not getting any stronger!
 16. I have an old tie that I can wear to the interview. My grandfather gave it to me last year.
 17. I had fun last night - we had quite a few drinks. But I have a really bad hangover this morning and I am considering calling in sick.
 18. Please don't interrupt me when I am talking! Now I have forgotten what I wanted to tell you.
 19. It was such a pleasure talking to you. I hope we stay in touch.
 20. I can't believe I misplaced my keys yet again! I have to leave for the airport right now.
 21. Gosh! the boy jumped right off the cliff into the ocean. He is lucky he didn't hit a rock.
 22. The baby just spit up on my brand new clothes. I am going to be late for our dinner tonight.
 23. The food smells disgusting but tastes delicious. How strange is that!
 24. I was about to park when I saw a person with a gun. I kept driving and called the police right away.
 25. I decided to take a nap during my lunch break. I am so glad I did! I feel very refreshed.
 26. The food didn't get delivered on time. We had to keep our guests waiting while we searched for options.
 27. I was walking down an alley the other night. I had the strange feeling that someone was following me.
 28. She twisted her ankle while ice-skating. It was her final performance for the season.
 29. Who moved my boxes from this room? I need to find my shoes before I can head out.
 30. We miss our old home in the mountains quite a bit. This new place just doesn't feel as cozy.

Subject demographics. Out of the total pool of subjects that volunteered data for our 2-stage data capture, 50% are female, 47.8% are male, and the remaining chose "a gender not listed here". Amongst different age groups, 37% of the participants are 25-34 years old, 32.6% are 35 – 44 years old, 17.4% are 45-54 years old, 6.5% are 18-24 years old, and 6.5% are 55-64 years old. In terms of race and ethnicity, 41.3% are Caucasian, 47.8% are Asian (including South Asian, East Asian, South-east Asian), 6.5% are African, 2.2% are Hispanic / Latino, 2.2% are Pacific islander, and others remained unspecified.

Synthetic Talking-Head Videos. As mentioned briefly in the main paper, we pool together videos for the 46 identities from our own 2-stage data capture, along with videos from 24 identities of RAVDESS (scripted mono-

Clip duration F	AUC
31	0.754
51	0.810
71	0.813
91	0.811

Table A3: Ablation study for different values of F.

logues only) [33], and 91 from CREMA-D (short scripted monologues only) [11], resulting in a total of 161 unique identities. For each of these 161 identities, the remaining 160 are used to drive cross-reenactments, with 8 driving videos randomly selected from the total set of videos for each driving identity. Moreover, all videos for the identity are used to drive self-reenactments. This results in at least 1300 synthetic videos per identity, with 1280 obtained from cross-reenactments. For any given target identity, we incorporate synthetic videos driven by *every* remaining identity. During training, such a large variety of cross-reenactments enable effectively learning an appearance-agnostic dynamic facial identity feature space.

A10. Implementation Details

The temporal ID net [14] is trained using our input features after appropriately modifying the number of input channels to match our feature dimension. To adjust the receptive field of the temporal ID net so that it predicts an embedding vector for longer or shorter input clips, we modify the number of layers of the network, and the dilation factor for the layers. Specifically, here are the kernel sizes and dilation factors for each of the layers in the temporal ID net, depending on the choice of input clip duration F:

1. F = 31 frames: (1, 1, 1, 1, 2, 2, 2, 2, 4)
2. F = 51 frames: (1, 1, 1, 1, 2, 2, 2, 4, 4, 4, 4)
3. F = 71 frames: (1, 1, 1, 1, 2, 2, 2, 2, 2, 2, 4, 4, 4, 4, 4)
4. F = 91 frames: (1, 1, 1, 1, 2, 2, 2, 2, 2, 2, 2, 2, 4, 4, 4, 4, 4, 4, 4)

The kernel sizes are all set to 3 apart from the first layer, which is 1. All other details of the temporal ID network are adapted from the existing implementation [14]. To implement the push and pull terms in Equations 1,2 in the main paper, n and t span over 5 consecutive F-frame clips in a video. That is, during training, the temporal ID net receives as input (F+4)-frame videos, and outputs 5 embedding vectors, one for each of the 5 F-frame clips in the video. The max operation in Equations 1,2 is performed over the 5 different clips (therefore, 5 different values of n), and the overall loss term in Equation 4 accumulates over 5 values of t . So, when a batch of videos is loaded for a training iteration, it comprises of (F+4)-frame "videos", which are split into 5 clips. These (F+4) frames are randomly selected from the entire video.

A11. Evaluation

In Table A3, we report the results of our experiment with varying values of F , which is the number of frames provided to the network to make a prediction about the dynamic facial temporal identity signature. We observe that increasing clip duration beyond 51 frames shows comparable performance to the result with $F=51$ (with a little AUC gain of about 0.003). Shorter clips might allow for computational efficiency and the ability to deploy our system for continuous online verification, but they also need to be long enough to capture meaningful expressions. Based on this experiment, we chose $F=51$ since it provides the best trade-off between the two competing requirements.

In Figure A6, we show more samples similar to the ones shown in Figure 5 of the main paper. For each row of results in Figure A6, we choose a reference identity, and a held-out set of reference self-reenacted videos for each of these identities. Then, we report the average Euclidean distance of the following videos with respect to the held-out self-reenacted videos for the reference identity:

1. a new self-reenacted video by the reference identity (not included in the held-out reference set) – highlighted with a green border,
2. a cross-reenacted video where the reference identity is the driver – highlighted with a green border, and
3. two cross-reenacted videos where the reference identity is the target, driven by some other identity – highlighted with a red and a blue border.

Based on the reported distance values, we observe that videos where the reference identity is the driver are closer to the set of other self-reenacted videos driven by the reference identity and far from those where reference identity is the target to be driven by other identities. This further confirms the ability of our model to fingerprint synthetic avatars based purely on facial motion, independent of the appearance of a synthetic talking-head video.




Figure A6: Animated figure. Open in a media-enabled viewer like Adobe Reader and click on the inset. Continuing Figure 5 from the main paper, we show more visual results to demonstrate that our method indeed predicts embedding vectors that lie close together when the clips have the same driving identity. As a reminder, for each row, we pick a reference identity. The **green box** indicates reenactments driven by the reference identity, the **red** and **blue** are cross-reenactments of the reference identity. We compute the average distance of each clip shown here against all other clips driven by the reference identity. The average distance to the other clips of the reference identity is consistent for a given motion, and lower (better) when the reference identity is driving as compared to the cross-reenactments that use the reference identity as target. Note that the indexing used to denote an identity (such as “N” in ID_N) is only intended to convey when an identity is different from another. These indices are not global and what is denoted as ID₁ in this figure may not match another figure.