

# That's What I Said: Fully-Controllable Talking Face Generation

Youngjoon Jang\*  
wgs01088@kaist.ac.kr  
KAIST  
Daejeon, Republic of Korea

Kyeongha Rho\*  
khrho325@kaist.ac.kr  
KAIST  
Daejeon, Republic of Korea

Jongbin Woo  
jongbin.woo@kaist.ac.kr  
KAIST  
Daejeon, Republic of Korea

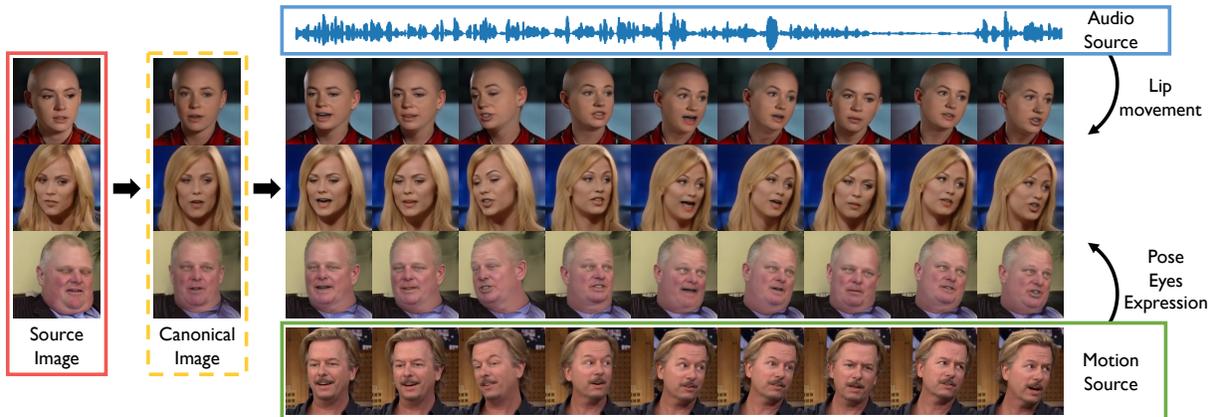
Hyeongkeun Lee  
lhk528@kaist.ac.kr  
KAIST  
Daejeon, Republic of Korea

Jihwan Park  
jihwan.park@42dot.ai  
42dot Inc., Hyundai Motor Company  
Seoul, Republic of Korea

Youshin Lim  
youshin.lim@42dot.ai  
42dot Inc., Hyundai Motor Company  
Seoul, Republic of Korea

Byeong-Yeol Kim  
byeongyeol.kim@42dot.ai  
42dot Inc., Hyundai Motor Company  
Seoul, Republic of Korea

Joon Son Chung  
joonson@kaist.ac.kr  
KAIST  
Daejeon, Republic of Korea



**Figure 1: Our Fully-Controllable Talking Face Generation (FC-TFG) framework precisely reflects every facial expression of the motion source while synchronising the lip shape with the input audio source. The key to our framework is to find the canonical space, where every face has the same motion patterns, but has different identities.**

## ABSTRACT

The goal of this paper is to synthesise talking faces with controllable facial motions. To achieve this goal, we propose two key ideas. The first is to establish a canonical space where every face has the same motion patterns but different identities. The second is to navigate a multimodal motion space that only represents motion-related features while eliminating identity information. To disentangle

identity and motion, we introduce an orthogonality constraint between the two different latent spaces. From this, our method can generate natural-looking talking faces with fully controllable facial attributes and accurate lip synchronisation. Extensive experiments demonstrate that our method achieves state-of-the-art results in terms of both visual quality and lip-sync score. To the best of our knowledge, we are the first to develop a talking face generation framework that can accurately manifest full target facial motions including lip, head pose, and eye movements in the generated video without any additional supervision beyond RGB video with audio.

\* These authors contributed equally.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

MM '23, October 29–November 3, 2023, Ottawa, ON, Canada.

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0108-5/23/10.

<https://doi.org/10.1145/3581783.3612587>

## CCS CONCEPTS

• Information systems → Multimedia content creation; • Computing methodologies → Image processing.

## KEYWORDS

Talking Face Generation, Face Animation, Motion Transfer

**ACM Reference Format:**

Youngjoon Jang\*, Kyeongha Rho\*, Jongbin Woo, Hyeongkeun Lee, Jihwan Park, Youshin Lim, Byeong-Yeol Kim, and Joon Son Chung. 2023. That's What I Said: Fully-Controllable Talking Face Generation. In *Proceedings of the 31st ACM International Conference on Multimedia (MM '23), October 29–November 3, 2023, Ottawa, ON, Canada*. ACM, Ottawa, Canada, 10 pages. <https://doi.org/10.1145/3581783.3612587>

**1 INTRODUCTION**

Audio-driven talking face generation technology has numerous applications in the film and entertainment industry, including virtual assistants, video conferencing, and dubbing. Its primary goal is to generate animated faces that closely match the audio, creating more engaging and interactive experiences for users. This technology has the potential to revolutionise various industries by making human-machine interactions more natural and immersive.

Previous literature on deep learning-based talking face synthesis can be divided into two branches. The first branch [7, 9, 30, 33, 43, 47, 61, 64] only uses RGB modality as a form of supervision to reconstruct a target video, while the second branch [2, 6, 8, 12, 22, 34, 42, 44, 45, 53, 63] leverages 2D or 3D structural information to propagate more detailed supervision. These works have made significant progress in generating natural-looking lip motions. Nevertheless, there is still a need for further development in controlling head pose and refining facial attributes in greater detail. Some methods [19, 27, 29, 50, 62] have the ability to generate talking faces that closely resemble the movements and identities of a target video. However, these methods still have limitations that hinder their practical application. For instance, [62] is only capable of altering the head pose, while [29] is incapable of creating detailed facial components like eye gaze movements. [19, 27, 50] require the utilisation of facial keypoints to physically separate facial components such as eyes and lips to create varied expressions.

In this paper, we propose a novel framework, Fully-Controllable Talking Face Generation (FC-TFG), which aims to generate talking faces that exactly copy full target motion including head pose, eyebrows, eye blinks, and eye gaze movements. The key to our method is preserving semantically meaningful features associated with the identity of each person via navigation of the latent space without requiring additional supervision such as facial keypoints.

Our work is raised from the fundamental question in talking face generation: is it possible to completely disentangle facial motions and facial identities in the latent space? Recent works [17, 21, 36, 48] discover that semantically meaningful directions exist in the latent space of Generative Adversarial Networks (GANs) by generating face images with target facial attributes. Based on these findings, [55] proposes a method that uses pre-defined orthogonal vectors representing basic visual transformations needed for face animations. However, using the motion dictionary fails to manipulate every detail of facial attributes such as eye gaze movements and lip shape due to the limited granularity of motion vectors.

Unlike the previous methods, we grant a higher degree of freedom to the motion vectors rather than restricting them to be constructed by a few orthogonal vectors. To achieve this goal, we disentangle the latent space of StyleGAN into two distinct subspaces (as shown in Fig. 2): (1) a canonical space that can accommodate

a variety of facial identities while maintaining consistent facial attributes, and (2) a multimodal motion space that contains exclusive motion features acquired by fusing both audio and image sources for transferring the target's motion to the source face image. We enforce that the multimodal motion codes are entirely disentangled from the canonical features by introducing an orthogonality constraint between the canonical space and the multimodal motion space. As a result, the proposed FC-TFG framework manipulates the latent code with a straightforward linear operation, which enables our system to generate more controllable facial animations while avoiding the unwanted entanglement of different features.

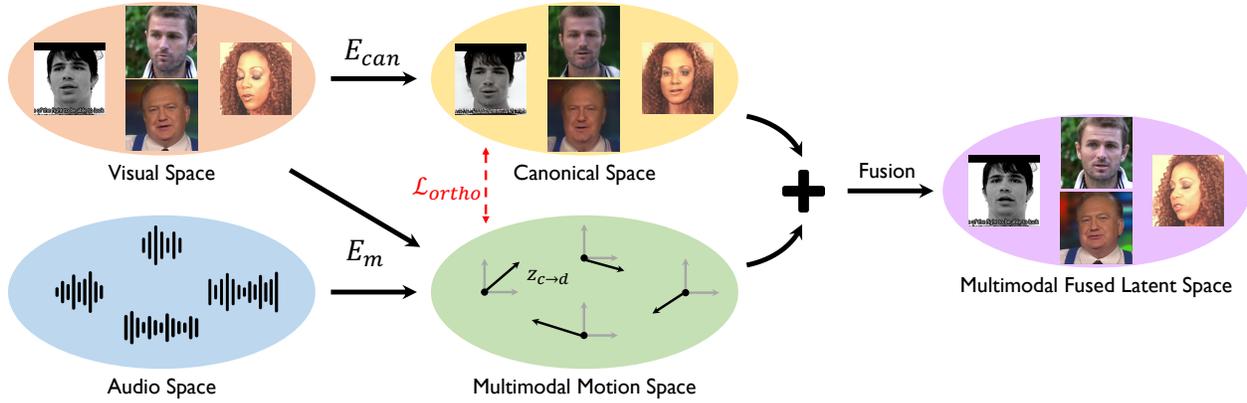
Furthermore, our framework is designed as a single encoder-decoder network that utilises StyleGAN's inversion network as the encoder. By taking advantage of this architecture, FC-TFG can effectively disentangle input images into canonical and motion features using simple MLP layers added on top of the inversion network, without significantly increasing the model complexity.

In summary, we make three key contributions: (1) We propose a novel framework, Fully-Controllable Talking Face Generation (FC-TFG), that generates talking faces with controllable target motion, including head pose, eyebrows, eye blinks, eye gaze, and lip movements. (2) We separate the style latent space into a canonical space that contains only person-specific characteristics and a multimodal motion space that contains person-agnostic motion features encoded from driving pose video and audio source. By imposing an orthogonality constraint on the correlation between the two spaces, the proposed model produces detailed and controllable facial animation. (3) We demonstrate that the proposed FC-TFG is highly effective in generating talking faces with sophisticated motion control, producing state-of-the-art results in both qualitative and quantitative metrics. This success highlights the potential of FC-TFG for a wide range of applications that demand elaborate control over various facial features.

**2 RELATED WORKS****2.1 Audio-Driven Talking Face Generation.**

The synthesis of speech-synchronised video portraits has been a significant area of research in computer vision and graphics for many years [5, 65]. Early works [15, 16] focus on the individual speaker setting where a single model can generate various talking faces corresponding to a single identity. Recently, the development of deep learning has enabled the design of more generalised talking face generation models [9, 26, 30, 33, 43, 61, 63] that produce talking faces by inputting identity conditions. However, these works neglect head movements due to the challenge of disentangling head poses from identity-related facial characteristics.

To produce talking face videos that have dynamic and natural movements, there exist studies that utilise landmarks or mesh information [8, 12, 42, 45, 59]. For the similar purpose, several works [3, 14, 23, 28, 53] adopt 3D information as intermediate representations. Even with additional modalities, the proposed models have certain limitations such as a lack of pose control and the inability to render faces with visually pleasing quality under one-shot conditions. Additionally, they suffer from severe performance degradation in situations where the accuracy of landmarks is low, especially in the wild.



**Figure 2: Overall framework.** Our goal is to transfer the motions of the individual portrait via latent space exploration. The overall pipeline is comprised of two steps. The first step involves transferring the code in the style latent space to the canonical latent space, where each face has the same facial attributes, but distinct identities. In the second step, the multimodal motion space is navigated, where each code contains only motion information and excludes identity information. The target motion code is obtained by fusing image latent features and audio latent features. The canonical code and the motion code are then combined using a linear operation, resulting in the final multimodal fused latent code. This fused latent code is used as input to the decoder to synthesise attribute-controllable talking face videos.

Several recent works [4, 19, 27, 29, 50, 62] have demonstrated that it is possible to create realistic talking faces that mimic different movements and identities from a target video. However, there are a number of limitations: [62] can only modify head pose, [29] cannot animate eye gazes, and [19, 27, 50] require facial keypoints to separate visual information for generating detailed expressions.

Unlike the previous works, our framework can generate a wide range of target facial features including pose, lip, eye blink, and even eye gaze without requiring extra annotations or structural information. Our framework seeks a canonical space of a generator, where each face has the same lip shape and pose but different identities. This method enables the creation of more advanced facial representations and simplifies the modeling of motion transfer between source and target images by changing their motion relationship from relative to absolute. Additionally, our model consistently produces high-quality videos under one-shot conditions.

## 2.2 Latent Space Editing.

Latent space editing involves intentionally modifying the generated output results by exploring meaningful directions in a high-dimensional latent space of a pre-trained generator network. These directions enable intuitive navigation that corresponds to desired image manipulation. To manipulate the latent space, some approaches [17, 20, 37] directly propagate labeled supervision such as facial attributes. Other works [32, 38, 49, 58] demonstrate the potential of modifying the semantics of images without annotation.

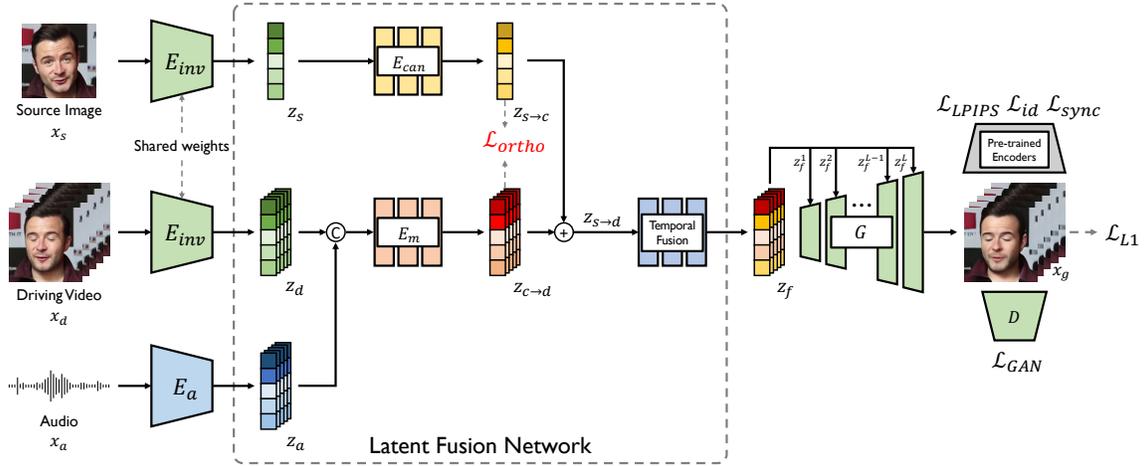
A recent work [55] achieves success in applying latent space editing technology to face reenactment task which aims to transfer target motion to source image using RGB modality alone.

In contrast to finding directions that correspond to individual facial movements in the latent space by using a single modality, our work aims to disentangle the face identity and the complex attributes composing face motions with multimodal features representing both visual and audio information.

## 3 METHOD

In this work, we propose a self-supervised approach for generating realistic talking face videos by transferring complex motions such as lip movements, head poses, eye gazes from a driving video to a source identity image. Such motion transformation is modeled via disentangled latent feature manipulation. The proposed framework’s pipeline is illustrated in Fig. 2. The key to our framework is two mapping operations: (1) Visual Space to Canonical Space, (2) Visual/Audio Space to Multimodal Motion Space. Through the first mapping, we can obtain canonical images that have the same motion features and different identities. Meanwhile, the second mapping yields motion vectors that enable us to transfer desired motions onto canonical images. To ensure the disentanglement of the two subspaces, we impose an orthogonality constraint. Based on this process, our model is capable of generating natural-looking talking face videos that mimic the full facial motions of the target.

We further provide an overall architecture of the proposed framework in Fig. 3. Our framework includes a visual encoder  $E_{inv}$ , an audio encoder  $E_a$ , a generator  $G$ , a discriminator  $D$ , and a Latent Fusion Network. We adopt a single shared visual encoder to extract both source and driving latent codes, denoted as  $z_s$  and  $z_d$ , respectively. The audio encoder extracts an additional audio latent code,  $z_a$ . The canonical encoder  $E_{can}$  maps  $z_s$  to a canonical code, denoted as  $z_{s→c}$ , which is then combined with a motion code  $z_{c→d}$  obtained from the multimodal motion encoder  $E_m$ . Since an orthogonality constraint is imposed between  $z_{s→c}$  and  $z_{c→d}$ , the motion transfer process is implemented by simply adding the two latent codes. To ensure generated motions are natural, the fused code  $z_{s→d}$  is passed through a Temporal Fusion layer before being fed into the generator  $G$ . The generator, which is the StyleGAN2 decoder, allows for controlling the coarse-fine motion transfers. Overall, the proposed method presents a promising approach for generating high-quality talking face videos with controllability.



**Figure 3: Overall Architecture of FC-TFG.** Our model consists of several components, including a visual encoder  $E_{inv}$ , an audio encoder  $E_a$ , a generator  $G$ , a discriminator  $D$ , and Latent Fusion Network. For compactness, we use a single visual encoder to extract both a source latent code  $z_s$  and a driving latent code  $z_d$ . The audio encoder extracts an audio latent code  $z_a$ . First, we design a canonical encoder  $E_{can}$  to map  $z_s$  to a canonical space as  $z_s \rightarrow z_{s \rightarrow c}$ , and then linearly combine it with a target motion code  $z_{c \rightarrow d}$ . This motion code is produced by a multimodal motion encoder  $E_m$ , which combines  $z_d$  and  $z_a$ . To generate natural motions, we pass the target code through a Temporal Fusion layer before feeding it into the generator  $G$ . The generated video  $x_g$  is compared to a driving video  $x_d$  based on its visual and synchronisation quality. Note that we employ the StyleGAN2 generator as a decoder to control coarse-fine motion transfers, and  $L$  is the number of modulation layers of the generator. Therefore, each model ( $E_{can}$  and  $E_m$ ) in Latent Fusion Network has  $L$  independent weights.

### 3.1 Navigating Canonical Space

Our main target is to obtain a latent code  $z_{s \rightarrow d; t} \sim \mathcal{Z} \in \mathbb{R}^N$  that captures the intricate motion transformation from the source image  $x_s$  to the  $t$ -th driving frame  $x_{d; t}$ . From now on, we omit the temporal index  $t$  for better readability. Directly seeking  $z_{s \rightarrow d}$  in the latent space is an arduous task because the model must be able to capture the subtle relative motion relationship between  $x_s$  and  $x_d$  while accurately representing the intricate facial attributes [55]. To address this challenging issue, we adopt an innovative approach proposed in [39, 40, 54] by assuming the existence of a canonical image  $x_c$  in a canonical space that has unified face-related motions but individual identities. Consequently, we can acquire the target mapping code  $z_{s \rightarrow d}$  through a two-stage motion transfer process.

$$z_{s \rightarrow d} = z_{s \rightarrow c} + z_{c \rightarrow d}, \quad (1)$$

where  $z_{s \rightarrow c}$  and  $z_{c \rightarrow d}$  indicate the transformation  $x_s \rightarrow x_c$  and  $x_c \rightarrow x_d$  respectively. Now, the mapping  $z_{s \rightarrow d}$  can be solved with an absolute motion transfer process, where both  $z_{s \rightarrow c}$  and  $z_{c \rightarrow d}$  do not need to account for the motion relationship between them.

Unlike the aforementioned previous works [39, 40, 54], which rely solely on the RGB modality to discover the canonical space, our method leverages both RGB and audio modalities and independently control the facial motions and lip movements simultaneously.

To design this motion transfer process effectively and intuitively, we apply a simple linear operation that adds  $z_{c \rightarrow d}$  directly to  $z_{s \rightarrow c}$ . Here,  $z_{s \rightarrow c}$  is obtained by passing  $z_s = E_{inv}(x_s)$  through  $E_{can}$  which is a 2-layer MLP, and  $z_{c \rightarrow d}$  is learned by navigating motion space that will be further explained in the following section.

### 3.2 Navigating Multimodal Motion Space

In order to generate natural head poses, lip motions, and facial expressions simultaneously,  $z_{c \rightarrow d}$  should successfully combine both vision and audio information while only containing motion-related features. To disentangle identity-related codes and motion-related codes in  $N$ -dimensional latent space, we inject a channel-wise orthogonality constraint into each channel in  $z_{s \rightarrow c}$  and  $z_{c \rightarrow d}$ .

By strictly enforcing this constraint, the model induces that the identity and motion codes do not interfere with each other when they are combined for the transformation with the linear operation as explained in Sec. 3.1. Finally,  $z_{c \rightarrow d}$  can be obtained as follows:

$$z_{c \rightarrow d} = E_m(E_a(x_a) \oplus E_{inv}(x_d)), \quad (2)$$

where  $E_a$ ,  $E_m$ , and  $x_a$  denote audio encoder, motion encoder, and audio input respectively.  $\oplus$  operation indicates channel-wise concatenation. Note that as we solely traverse the latent space of the pre-trained image decoder, the resulting  $z_{s \rightarrow d}$  can retain semantically meaningful features. The structure of the motion encoder is a 3-layer MLP. Finally, to generate temporally consistent latent features, we refine the acquired  $z_{s \rightarrow d}$  by feeding it to the Temporal Fusion layer, which consists of a single 1D convolutional layer.

### 3.3 Training Objectives

**3.3.1 Orthogonality loss.** To ensure effective disentangling of  $z_{s \rightarrow c}$  and  $z_{c \rightarrow d}$ , we extend the application of the orthogonality loss [58] to each layer of the StyleGAN decoder while combining both audio and video information. This approach enables us to manipulate finer movements, enhancing the overall capability of our framework. The orthogonality loss function is expressed as follows:

Method	Controllable Motions				Voxceleb2					MEAD				
	lip	pose	eye blink	eye gaze	SSIM↑	MS-SSIM↑	PSNR↑	LMD↓	LSE-C↑	SSIM↑	MS-SSIM↑	PSNR↑	LMD↓	LSE-C↑
Wav2Lip [33]	✓	✗	✗	✗	0.58	-	20.63	2.65	<b>8.66</b>	<b>0.85</b>	-	26.15	3.11	<b>7.25</b>
MakeItTalk [63]	✓	✗	✗	✗	0.55	0.45	16.94	3.28	3.83	0.78	0.78	23.6	3.55	4.35
Audio2Head [52]	✓	✗	✗	✗	0.51	0.40	15.98	3.58	5.79	0.67	0.59	19.57	4.85	5.38
PC-AVS [62]	✓	✓	✗	✗	0.57	0.60	17.37	2.25	5.82	0.66	0.7	19.95	2.93	5.19
FC-TFG (Ours)	✓	✓	✓	✓	<b>0.69</b>	<b>0.77</b>	<b>21.22</b>	<b>1.58</b>	8.46	0.84	<b>0.89</b>	<b>26.19</b>	<b>2.46</b>	5.51

**Table 1: Quantitative Results. We compare our method to four publicly available baselines on six different metrics. Our approach outperforms the baseline methods in terms of both visual quality and lip synchronisation, while simultaneously controlling diverse and detailed facial motions. We evaluate the generated samples using the original authors’ experimental settings, ensuring a fair comparison between the different methods.**

$$\mathcal{L}_{ortho} = \frac{1}{N} \sum (z_{s \rightarrow c} \odot z_{c \rightarrow d}), \quad (3)$$

where  $N$  denotes the number of layers in decoder and  $\odot$  indicates Hadamard product operation.

**3.3.2 Synchronisation loss.** To generate well-synchronised videos  $x_g$  according to the input audio conditions, we adopt a sync loss function, which leverages a pre-trained SyncNet [11] comprising an audio encoder and a video encoder. Many works [28–30, 33, 53] modify the SyncNet with altered objective functions, to further improve lip synchronisation quality. We use the modified SyncNet introduced in [29] to enhance our model’s lip representations.

The distance between the features of a video and its synchronised audio, extracted from the pre-trained Syncnet model, should be close to 0. From this, we minimise the following sync loss:

$$\mathcal{L}_{sync} = 1 - \cos(f_v(x_g), f_a(x_a)), \quad (4)$$

where  $f_a$  and  $f_v$  denote the audio encoder and video encoder of SyncNet respectively.

**3.3.3 Identity loss.** To preserve facial identity after motion transformation, we apply an identity-based similarity loss [35] by employing a pre-trained face recognition network  $E_{id}$  [13].

$$\mathcal{L}_{id} = 1 - \cos(E_{id}(x_g), E_{id}(x_d)). \quad (5)$$

**3.3.4 Reconstruction loss.** For the reconstruction loss, we adopt  $L1$  loss function that calculates the pixel-wise  $L1$  distance between the generated talking face image,  $x_g$ , and the target image,  $x_d$ . The reconstruction loss can be calculated as follows:

$$\mathcal{L}_{rec} = \|x_g - x_d\|_1. \quad (6)$$

**3.3.5 Perceptual loss.** Using  $L1$  reconstruction loss alone may result in blurry images or slight artifacts as it is a pixel-level loss. To compensate for the smoothing effect caused by the reconstruction loss  $\mathcal{L}_{L1}$ , we add the Learned Perceptual Image Patch Similarity (LPIPS) loss [60], which measures the perceptual similarity between two images. LPIPS loss can be calculated as follows:

$$\mathcal{L}_{LPIPS} = \frac{1}{N_f} \sum_{i=1}^{N_f} \|\phi(x_g)_i - \phi(x_d)_i\|_2, \quad (7)$$

where  $\phi$  is a pre-trained VGG19 [41] network, and  $N_f$  is the number of feature maps.

**3.3.6 Adversarial loss.** We perform adversarial training with an image discriminator,  $D$ , to improve the quality of the generated videos. The architecture of  $D$  is the same as StyleGAN2 [24] discriminator. We use a non-saturating loss [18] for adversarial training, which can be expressed as follows:

$$\mathcal{L}_{GAN} = \min_G \max_D \left( \mathbb{E}_{x_d} [\log(D(x_d))] + \mathbb{E}_{z_f} [\log(1 - D(G(z_f)))] \right). \quad (8)$$

**3.3.7 Overall loss.** Our total loss can be formulated as follows:

$$\mathcal{L}_{total} = \lambda_1 \mathcal{L}_{ortho} + \lambda_2 \mathcal{L}_{sync} + \lambda_3 \mathcal{L}_{id} + \lambda_4 \mathcal{L}_{rec} + \lambda_5 \mathcal{L}_{LPIPS} + \lambda_6 \mathcal{L}_{GAN}, \quad (9)$$

where hyperparameters  $\lambda$  are introduced to balance the scale of each loss. Each  $\lambda$  controls the relative importance of its corresponding loss term. Empirically,  $\lambda_1, \lambda_2, \lambda_3, \lambda_4, \lambda_5$ , and  $\lambda_6$  are set to 1, 0.1, 0.5, 1, 1 and 0.1 respectively.

## 4 EXPERIMENTS

### 4.1 Experimental Setup

**4.1.1 Dataset.** Our framework is trained on VoxCeleb2 [10] and evaluated on both VoxCeleb2 and MEAD [51]. VoxCeleb2 includes 6,112 different identities and over 1 million utterances. Of the total identities, 5,994 are used for training, while the remaining identities are reserved for testing. We follow the pre-processing procedure proposed in [39] to ensure the consistent visual quality of videos.

To evaluate one-shot talking face generation performance, we utilise the MEAD dataset, which comprises emotional faces featuring more than 30 actors and eight emotion categories at three intensity levels. To conduct the evaluation, we randomly choose five speakers and five videos for each emotion category, and we only use the frontal-view videos from this dataset for testing purposes.

**4.1.2 Implementation Details.** First of all, we pre-train a StyleGAN2 [24] generator on the VoxCeleb2 dataset and then train HyperStyle [1] inversion network with the pre-trained StyleGAN2 model. Specifically, we replace the e4e [46] encoder in the HyperStyle model with pSp [35] encoder. We focus on manipulating 8 specific layers of 14 layers in our generator, namely layers 1, 2, 3, 4, 7, 8, 9, and 10. Additionally, we only input the audio feature into 2 specific layers, layers 7 and 8. This allows us to effectively control the style of the generated images based on the audio input.

For the audio source, we downsample the audio to 16kHz, then convert the downsampled audio to mel-spectrograms with a window size of 800, a hop length of 200, and 80 Mel filter banks. We utilise a pre-trained audio encoder introduced in [33]. We use Adam [25] optimiser for updating our model, with a learning rate of  $1e-4$ . Our framework is implemented on PyTorch [31] and trained with eight 48GB A6000 GPUs. Note that the aforementioned pre-trained models are fine-tuned during the training stage.

**4.1.3 Comparison Methods.** We compare our method with several state-of-the-art talking face synthesis methods, which baselines are currently available. **Wav2Lip** [33] employs pre-trained SyncNet as a lip-sync discriminator to generate well-synchronised mouth region of the source image. **MakeItTalk** [63] predicts landmarks through 3D face models and generates both lip movements and head motions simultaneously driven by audio. **Audio2Head** [52] generates head motions by utilising a keypoint-based dense motion field driven by audio. **PC-AVS** [62] is a pose-controllable talking face generation model that controls head poses with driving videos.

## 4.2 Quantitative Results

**4.2.1 Evaluation Metrics.** We conduct quantitative evaluations with various evaluation metrics that have previously been adopted in the talking face generation field. To account for the accuracy of mouth shapes and lip sync, we use Landmarks Distance (LMD) around the mouths proposed in [8], and Lip Sync Error Confidence (LSE-C) proposed in [11]. To compare the visual quality of the generated videos, we calculate Structural Similarity Index Measure (SSIM) [56], Multi Scale Structural Similarity Index Measure (MS-SSIM) [57], and Peak Signal-to-Noise Ratio (PSNR).

**4.2.2 Controllable Motion Types.** In order to assess the capabilities of the baseline methods in generating realistic talking faces, we compare the number of controllable motions for each method under the target conditions in Table 1. Wav2Lip, MakeItTalk, and Audio2Head rely solely on audio input to generate videos, resulting in limited control over motion and pose. In particular, MakeItTalk, and Audio2Head are capable of generating lip movements and small head pose variations that are synchronised with the audio conditions but are unable to produce more complex and diverse motion patterns. On the other hand, PC-AVS and Ours have the capacity to control head pose following the target pose conditions. However, while PC-AVS is limited to controlling head pose, our approach is effective for controlling all facial attributes, including eye blinks and gazes, by exploring disentangled latent spaces.

**4.2.3 Evaluation Results.** We follow the evaluation protocol introduced in [62]. Specifically, we select the first frame of each test video to serve as a reference identity. We use the remaining frames to determine the subject’s pose, facial expression, and lip shape.

As shown in Table 1, we compare our method with four baselines on VoxCeleb2 and MEAD datasets. For VoxCeleb2 dataset, except for LSE-C metric, our proposed method surpasses the previous audio-driven and pose-controllable methods in all metrics while manipulating detailed facial components. For MEAD dataset, our framework shows the best performance on MS-SSIM, PSNR, and LMD and comparable results to the Wav2Lip model in terms of SSIM. Although the Wav2Lip model shows better performance

on LSE-C metric, the generated face shows significant distortions. These results demonstrate that our model is capable to capture detailed facial movements and generate visually pleasing results by keeping semantically meaningful the latent space.

## 4.3 Qualitative Results

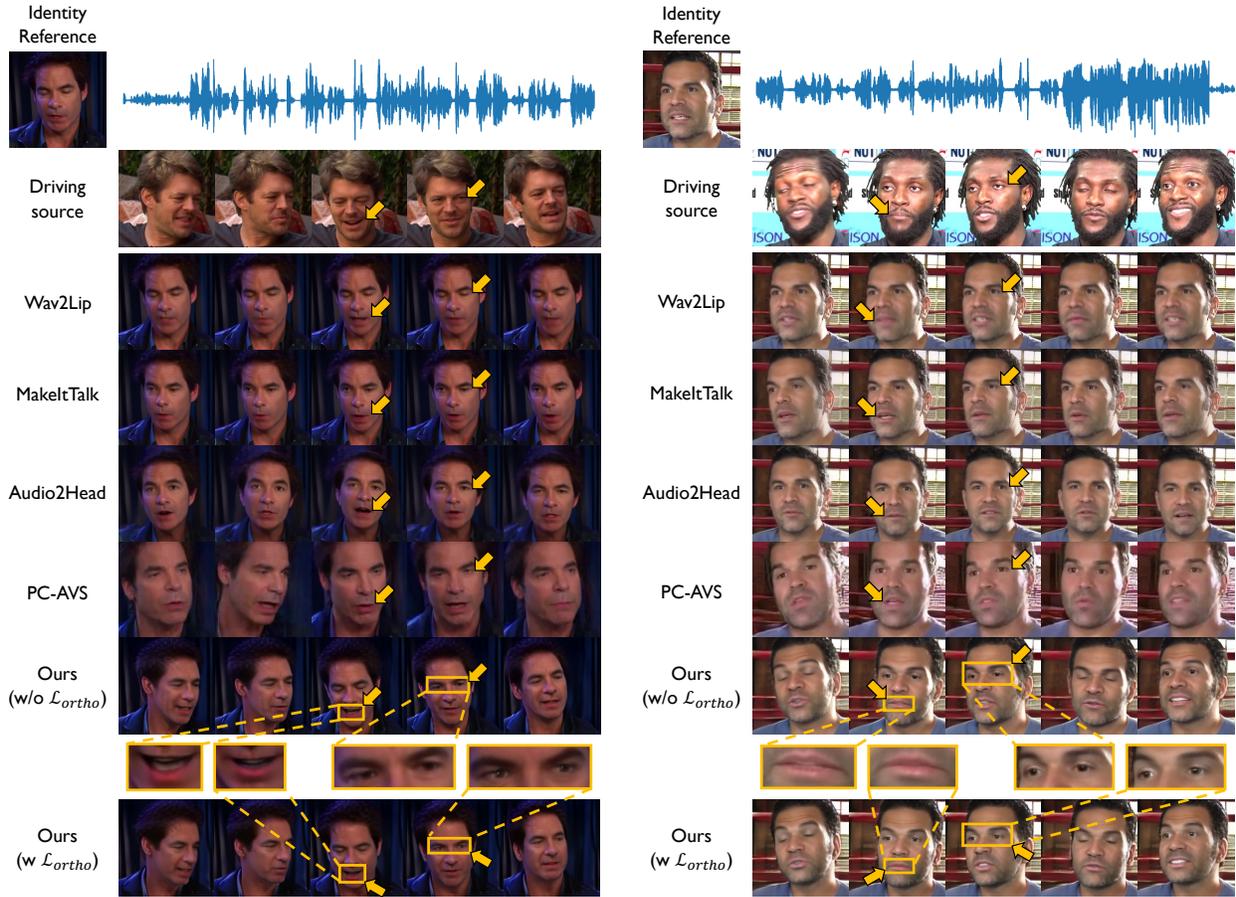
**4.3.1 Fully-Controllable Talking Face Generation.** We visually show our qualitative results in Fig. 4. We would like to clarify that Wav2Lip is a model that targets to change only the lip region. MakeItTalk and Audio2Head are models that generate natural face motions conditioned by only audio sources. For that reason, the generated videos cannot replicate diverse facial motions in the driving videos. PC-AVS focuses on synthesising the talking faces with controllable head pose variations. However, it fails to mimic the driving source’s facial expressions such as eye blinks, eyebrows movements, and eye gazes. On the other hand, our method can generate fully-controllable talking faces by synthesising both head pose and facial expressions that precisely follow the driving source videos.

In the left example of Fig. 4, our model captures the driving source’s smiles and eye gaze, as indicated by the yellow arrows. Similarly, in the right example, our model replicates detailed facial movements such as lip shape, eyebrow position, forehead frowns, and eye gazes. This high level of accuracy and detail allows our method to generate realistic and expressive talking faces that closely mimic the emotions and expressions of the driving source.

**4.3.2 User Study.** We assess the quality of the videos generated by FC-TFG with a user study of 40 participants for their opinions on 20 videos. Specifically, we randomly collect reference images, driving videos, and driving audios from VoxCeleb2 [10] test split. Subsequently, we create videos based on Wav2Lip [33], MakeItTalk [63], Audio2Head [52], PC-AVS [62], and FC-TFG (Ours). We adopt the widely used Mean Opinion Scores (MOS) as an evaluation metric following [27, 62]. Each user gives evaluation scores from 1 to 5 for the following aspects: (1) lip sync quality; (2) head movement naturalness; and (3) video realness. The order of methods within each video clip is randomly shuffled. Note that for a fair comparison with fully audio-driven generation methods, we do not evaluate how well the generated head poses are matched to the poses of the driving videos. As shown in Table 2, it is demonstrated that our method generates talking face videos with higher lip synchronisation and natural head movement, compared to the existing methods.

## 4.4 Ablation Study

**4.4.1 Canonical Space.** In Fig. 5 (a), we present the impressive outcomes of mapping various identities to the canonical space *where each face has the same facial motion but different identities*. The effectiveness of our approach in preserving identity can be clearly observed through the distinct and recognisable facial features of each individual. Additionally, as shown in Fig. 5 (b), regardless of the differences in the input images having the same identity, our approach consistently produces similar canonical images that accurately capture the individual’s distinct facial identity and detailed features. These results further highlight the effectiveness of our method in finding the canonical space that is essential for transferring complex facial motions with a simple linear operation.



**Figure 4: Qualitative Results.** We compare our method with several baselines listed in Table 1. Our approach outperforms all the baselines in terms of generating precise head motion and facial expressions following the given conditions. Wave2Lip, MakeItTalk, and Audio2Head fail to generate accurate head motion of the driving source videos. PC-AVS produces a similar head pose with target motion but lacks in generating realistic facial expressions. On the other hand, our method successfully generates every facial expression of the target motion while synchronising the lip with the input audio source.

Method	Wav2Lip	MakeItTalk	Audio2Head	PC-AVS	FC-TFG(Ours)
Lip Sync Quality↑	3.47	2.31	2.55	3.29	<b>3.93</b>
Head Movement Naturalness↑	1.88	2.5	2.98	3.23	<b>4.14</b>
Overall Quality↑	2.18	2.64	2.91	2.91	<b>3.94</b>

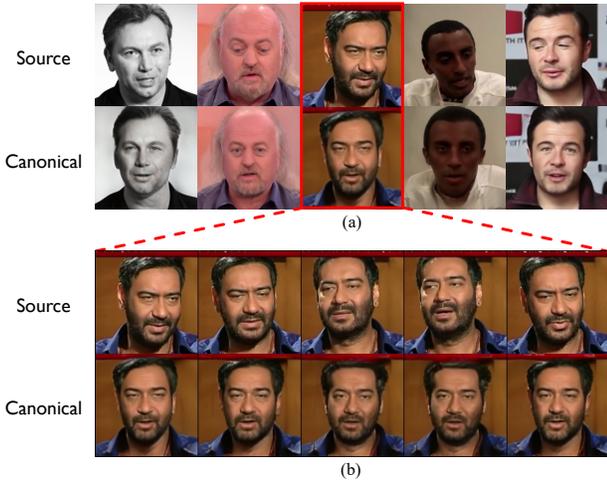
**Table 2: User Study.** We conduct a user study on generated videos with three aspects: lip synchronisation, the naturalness of head movement, and overall video quality. The higher the better, with the value range of 1 to 5.

**4.4.2 Ablations on Model Design.** We conduct ablation studies on our model choices for the Temporal Fusion layer and Discriminator  $D$  in Table 3. We empirically observe that a flickering phenomenon appears when using Long Short-Term Memory (LSTM) network as our temporal fusion model. We suspect that this issue is caused by our window size, which does not have any overlap along the temporal dimension for faster inference. We solve the flickering issue by replacing the LSTM layer with a 1D convolutional layer.

We investigate two different types of discriminators in this work; (1) image discriminator, and (2) video discriminator. The former is

designed to evaluate the authenticity of a single image input, while the latter is trained to classify a sequence of frames concatenated along the channel dimension as either real or fake video.

As shown in Table 3, the lip sync confidence score is high when we use the image discriminator. We believe that this is because the image discriminator is designed to analyse the details of a single image, which means that the generator must focus on spatial information to deceive the discriminator. On the other hand, the video discriminator examines the authenticity of multiple frames at the same time, which requires the generator to focus on both spatial



**Figure 5: Samples in Canonical Space.** We demonstrate how well our model preserves identity by mapping various identities to the canonical space. In Fig. (a), we generate diverse canonical image samples having different identities by feeding each canonical code to our generator. In Fig. (b), we further visualise every canonical image from a single video. These results prove that our model is robust to maintain the source identities and well-generalised to various identities.

Temporal Fusion		$D$		Metrics				
1D-Conv	LSTM	image	video	SSIM $\uparrow$	MS-SSIM $\uparrow$	PSNR $\uparrow$	LMD $\downarrow$	LSE-C $\uparrow$
	✓	✓		0.66	0.75	20.78	1.79	<b>8.66</b>
	✓		✓	0.67	0.75	20.64	1.73	7.9
✓			✓	0.66	0.75	20.72	1.98	8.04
✓		✓		<b>0.69</b>	<b>0.77</b>	<b>21.22</b>	<b>1.58</b>	<b>8.46</b>

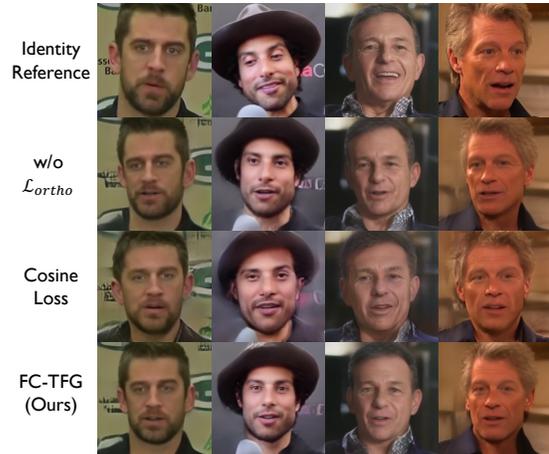
**Table 3: Ablations on Model Design.** We report performance results based on our different model choices. We observe that the highest performance is achieved when using a 1D convolutional layer for the Temporal Fusion layer and an image discriminator for the discriminator.

and temporal information. Based on our analysis, we conclude that the image discriminator is more suitable for our intended purpose of accurately generating lip shape, which is a spatially small region.

**4.4.3 Orthogonality Constraint Effectiveness.** We provide further insight into the efficacy of the orthogonality constraint by visually analysing samples in the canonical space. As depicted in Fig. 6, the canonical space of the model trained without orthogonality constraint contains diverse head poses, indicating the potential lack of disentanglement between identity and motion information. Additionally, when cosine similarity loss is used as the orthogonality constraint, the model produces even more blurry images. On the contrary, the canonical space of FC-TFG trained with  $L_{ortho}$  in Eqn. 3 contains more unified facial motions, indicating a successful disentanglement of identity and motion information in the latent space. The quantitative results according to the different orthogonality constraints are reported in Table 4. These results

Orthogonal Constraint	Metrics				
	SSIM $\uparrow$	MS-SSIM $\uparrow$	PSNR $\uparrow$	LMD $\downarrow$	LSE-C $\uparrow$
$\times$	0.67	0.76	20.88	1.70	8.27
Cosine	0.67	0.75	21	2.11	8.38
$L_{ortho}$	<b>0.69</b>	<b>0.77</b>	<b>21.22</b>	<b>1.58</b>	<b>8.46</b>

**Table 4: Ablations on Orthogonality Loss.** We ablate the effectiveness of the orthogonality constraint. The model trained with the proposed loss  $L_{ortho}$  shows the best performance.



**Figure 6: Analysis on the effectiveness of Orthogonality Constraint.** Our model effectively disentangles the identity and motion information by enforcing orthogonality between the canonical space and the multimodal motion space. This is demonstrated by the cohesive and consistent motion patterns in the canonical space, which is not achieved in the models either trained without the orthogonality constraint or trained with other kinds of constraint such as cosine loss.

strongly suggest that the orthogonality constraint plays a crucial role in achieving high-quality outcomes.

## 5 CONCLUSION

In this work, we propose a framework named Fully-Controllable Talking Face Generation (FC-TFG), which can generate every facial expression while synchronising lip movements with the input audio sources. Our framework is carefully designed to disentangle the latent space of StyleGAN into the canonical space and the multimodal motion space. This disentanglement enables the generation of more detailed and controllable facial animations while avoiding the unwanted mixing of different types of conditions. With various experiments, we prove that the proposed method is highly effective, achieving state-of-the-art results in both qualitative and quantitative evaluations of generated video quality. The potential of FC-TFG for a wide range of applications that demand precise control over various facial features is significant, such as virtual reality, and augmented reality. It could also be useful for creating personalised digital avatars or virtual assistants that can communicate and interact with users in a more natural and realistic manner.

## REFERENCES

- [1] Yuval Alaluf, Omer Tov, Ron Mokady, Rinon Gal, and Amit Bermano. 2022. Hyperstyle: Stylegan inversion with hypernetworks for real image editing. In *Proc. CVPR*. 18511–18521.
- [2] Robert Anderson, Björn Stenger, Vincent Wan, and Roberto Cipolla. 2013. An expressive text-driven 3d talking head. In *Proc. ACM SIGGRAPH*. 1–1.
- [3] Adrian Bulat and Georgios Tzimiropoulos. 2017. How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230,000 3d facial landmarks). In *Proc. ICCV*. 1021–1030.
- [4] Egor Burkov, Igor Pasechnik, Artur Grigorev, and Victor Lempitsky. 2020. Neural head reenactment with latent pose descriptors. In *Proc. CVPR*. 13786–13795.
- [5] Lele Chen, Guofeng Cui, Ziyi Kou, Haitian Zheng, and Chenliang Xu. 2020. What comprises a good talking-head video generation?: A survey and benchmark. *arXiv preprint arXiv:2005.03201* (2020).
- [6] Lele Chen, Guofeng Cui, Celong Liu, Zhong Li, Ziyi Kou, Yi Xu, and Chenliang Xu. 2020. Talking-head generation with rhythmic head motion. In *Proc. ECCV*. Springer, 35–51.
- [7] Lele Chen, Zhiheng Li, Ross K Maddox, Zhiyao Duan, and Chenliang Xu. 2018. Lip movements generation at a glance. In *Proc. ECCV*. 520–535.
- [8] Lele Chen, Ross K Maddox, Zhiyao Duan, and Chenliang Xu. 2019. Hierarchical cross-modal talking face generation with dynamic pixel-wise loss. In *Proc. CVPR*. 7832–7841.
- [9] Joon Son Chung, Amir Jamaludin, and Andrew Zisserman. 2017. You said that?. In *Proc. BMVC*.
- [10] Joon Son Chung, Arsha Nagrani, and Andrew Zisserman. 2018. VoxCeleb2: Deep speaker recognition. In *Proc. Interspeech*.
- [11] Joon Son Chung and Andrew Zisserman. 2017. Out of time: automated lip sync in the wild. In *Proc. ACCV*. Springer, 251–263.
- [12] Dipanjan Das, Sandika Biswas, Sanjana Sinha, and Brojeshwar Bhowmick. 2020. Speech-driven facial animation using cascaded gans for learning of motion and texture. In *Proc. ECCV*. Springer, 408–424.
- [13] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. 2019. Arcface: Additive angular margin loss for deep face recognition. In *Proc. CVPR*. 4690–4699.
- [14] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. 2019. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *Proc. CVPR*. 0–0.
- [15] Bo Fan, Lijuan Wang, Frank K Soong, and Lei Xie. 2015. Photo-real talking head with deep bidirectional LSTM. In *Proc. ICASSP*. IEEE, 4884–4888.
- [16] Bo Fan, Lei Xie, Shan Yang, Lijuan Wang, and Frank K Soong. 2016. A deep bidirectional LSTM approach for video-realistic talking head. *Multimedia Tools and Applications* 75 (2016), 5287–5309.
- [17] Lore Goetschalckx, Alex Andonian, Aude Oliva, and Phillip Isola. 2019. Analyze: Toward visual definitions of cognitive image properties. In *Proc. ICCV*. 5744–5753.
- [18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (2020), 139–144.
- [19] Geumbyeol Hwang, Sunwon Hong, Seunghyun Lee, Sungwoo Park, and Gyeongsu Chae. 2023. DisCoHead: Audio-and-Video-Driven Talking Head Generation by Disentangled Control of Head Pose and Facial Expressions. *arXiv preprint arXiv:2303.07697* (2023).
- [20] Ali Jahanian, Lucy Chai, and Phillip Isola. 2019. On the "steerability" of generative adversarial networks. *arXiv preprint arXiv:1907.07171* (2019).
- [21] Ali Jahanian, Lucy Chai, and Phillip Isola. 2020. On the "steerability" of generative adversarial networks. In *International Conference on Learning Representations*.
- [22] Xinya Ji, Hang Zhou, Kaisiyuan Wang, Wayne Wu, Chen Change Loy, Xun Cao, and Feng Xu. 2021. Audio-driven emotional video portraits. In *Proc. CVPR*. 14080–14089.
- [23] Zi-Hang Jiang, Qianyi Wu, Keyu Chen, and Juyong Zhang. 2019. Disentangled representation learning for 3D face shape. In *Proc. CVPR*. 11957–11966.
- [24] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2020. Analyzing and improving the image quality of stylegan. In *Proc. CVPR*. 8110–8119.
- [25] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [26] Prajwal KR, Rudrabha Mukhopadhyay, Jerin Philip, Abhishek Jha, Vinay Namboodiri, and CV Jawahar. 2019. Towards automatic face-to-face translation. In *Proc. ACM MM*. 1428–1436.
- [27] Borong Liang, Yan Pan, Zhizhi Guo, Hang Zhou, Zhibin Hong, Xiaoguang Han, Junyu Han, Jingtuo Liu, Errui Ding, and Jingdong Wang. 2022. Expressive talking head generation with granular audio-visual control. In *Proc. CVPR*. 3387–3396.
- [28] Yifeng Ma, Suzhen Wang, Zhipeng Hu, Changjie Fan, Tangjie Lv, Yu Ding, Zhi-dong Deng, and Xin Yu. 2023. StyleTalk: One-shot Talking Head Generation with Controllable Speaking Styles. *arXiv preprint arXiv:2301.01081* (2023).
- [29] Dongchan Min, Minyoung Song, and Sung Ju Hwang. 2022. StyleTalker: One-shot Style-based Audio-driven Talking Head Video Generation. *arXiv preprint arXiv:2208.10922* (2022).
- [30] Se Jin Park, Minsu Kim, Joanna Hong, Jeongsoo Choi, and Yong Man Ro. 2022. Synctalkface: Talking face generation with precise lip-syncing via audio-lip memory. In *Proc. AAAI*, Vol. 36. 2062–2070.
- [31] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library.
- [32] William Peebles, John Peebles, Jun-Yan Zhu, Alexei Efros, and Antonio Torralba. 2020. The hessian penalty: A weak prior for unsupervised disentanglement. In *Proc. ECCV*. Springer, 581–597.
- [33] KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Namboodiri, and CV Jawahar. 2020. A lip sync expert is all you need for speech to lip generation in the wild. In *Proc. ACM MM*. 484–492.
- [34] Alexander Richard, Colin Lea, Shugao Ma, Jurgen Gall, Fernando De la Torre, and Yaser Sheikh. 2021. Audio-and gaze-driven facial animation of codec avatars. In *Proc. WACV*. 41–50.
- [35] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. 2021. Encoding in style: a stylegan encoder for image-to-image translation. In *Proc. CVPR*. 2287–2296.
- [36] Yujun Shen, Jinjin Gu, Xiaou Tang, and Bolei Zhou. 2020. Interpreting the Latent Space of GANs for Semantic Face Editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- [37] Yujun Shen, Jinjin Gu, Xiaou Tang, and Bolei Zhou. 2020. Interpreting the latent space of gans for semantic face editing. In *Proc. CVPR*. 9243–9252.
- [38] Yujun Shen and Bolei Zhou. 2021. Closed-form factorization of latent semantics in gans. In *Proc. CVPR*. 1532–1540.
- [39] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. 2019. First order motion model for image animation. In *NeurIPS*, Vol. 32.
- [40] Aliaksandr Siarohin, Oliver J Woodford, Jian Ren, Menglei Chai, and Sergey Tulyakov. 2021. Motion representations for articulated animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13653–13662.
- [41] Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. In *Proc. ICLR*.
- [42] Linsen Song, Wayne Wu, Chen Qian, Ran He, and Chen Change Loy. 2022. Everybody's talkin': Let me talk as you want. *IEEE Transactions on Information Forensics and Security* 17 (2022), 585–598.
- [43] Yang Song, Jingwen Zhu, Dawei Li, Xiaolong Wang, and Hairong Qi. 2018. Talking face generation by conditional recurrent adversarial network. *arXiv preprint arXiv:1804.04786* (2018).
- [44] Supasorn Suwajanakorn, Steven M Seitz, and Ira Kemelmacher-Shlizerman. 2017. Synthesizing obama: learning lip sync from audio. *ACM Transactions on Graphics* 36, 4 (2017), 1–13.
- [45] Justus Thies, Mohamed Elgharib, Ayush Tewari, Christian Theobalt, and Matthias Nießner. 2020. Neural voice puppetry: Audio-driven facial reenactment. In *Proc. ECCV*. Springer, 716–731.
- [46] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. 2021. Designing an encoder for stylegan image manipulation. *ACM Transactions on Graphics* 40, 4 (2021), 1–14.
- [47] Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. 2020. Realistic speech-driven facial animation with gans. *Intl. Journal of computer vision* 128 (2020), 1398–1413.
- [48] Andrey Voynov and Artem Babenko. 2014. Unsupervised discovery of interpretable directions in the gan latent space. *arXiv preprint arXiv:2002.03754* (2014).
- [49] Andrey Voynov and Artem Babenko. 2020. Unsupervised discovery of interpretable directions in the gan latent space. In *Proc. ICML*. PMLR, 9786–9796.
- [50] Duomin Wang, Yu Deng, Zixin Yin, Heung-Yeung Shum, and Baoyuan Wang. 2022. Progressive Disentangled Representation Learning for Fine-Grained Controllable Talking Head Synthesis. *arXiv preprint arXiv:2211.14506* (2022).
- [51] Kaisiyuan Wang, Qianyi Wu, Linsen Song, Zhuoqian Yang, Wayne Wu, Chen Qian, Ran He, Yu Qiao, and Chen Change Loy. 2020. Mead: A large-scale audio-visual dataset for emotional talking-face generation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI*. Springer, 700–717.
- [52] Suzhen Wang, Lincheng Li, Yu Ding, Changjie Fan, and Xin Yu. 2021. Audio2Head: Audio-driven One-shot Talking-head Generation with Natural Head Motion. In *Proc. IJCAI*.
- [53] Suzhen Wang, Lincheng Li, Yu Ding, and Xin Yu. 2022. One-shot talking face generation from single-speaker audio-visual correlation learning. In *Proc. AAAI*, Vol. 36. 2531–2539.
- [54] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. 2021. One-shot free-view neural talking-head synthesis for video conferencing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10039–10049.
- [55] Yaohui Wang, Di Yang, Francois Bremond, and Antitza Dantcheva. 2022. Latent Image Animator: Learning to Animate Images via Latent Space Navigation. In *International Conference on Learning Representations*.
- [56] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612.

- [57] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. 2003. Multiscale structural similarity for image quality assessment. In *The Asilomar Conference on Signals, Systems & Computers*, Vol. 2. Ieee, 1398–1402.
- [58] Guoxing Yang, Nanyi Fei, Mingyu Ding, Guangzhen Liu, Zhiwu Lu, and Tao Xiang. 2021. L2m-gan: Learning to manipulate latent space semantics for facial attribute editing. In *Proc. CVPR*. 2951–2960.
- [59] Ran Yi, Zipeng Ye, Juyong Zhang, Hujun Bao, and Yong-Jin Liu. 2020. Audio-driven talking face video generation with learning-based personalized head pose. *arXiv preprint arXiv:2002.10137* (2020).
- [60] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. CVPR*. 586–595.
- [61] Hang Zhou, Yu Liu, Ziwei Liu, Ping Luo, and Xiaogang Wang. 2019. Talking face generation by adversarially disentangled audio-visual representation. In *Proc. AAAI*, Vol. 33. 9299–9306.
- [62] Hang Zhou, Yasheng Sun, Wayne Wu, Chen Change Loy, Xiaogang Wang, and Ziwei Liu. 2021. Pose-controllable talking face generation by implicitly modularized audio-visual representation. In *Proc. CVPR*. 4176–4186.
- [63] Yang Zhou, Xintong Han, Eli Shechtman, Jose Echevarria, Evangelos Kalogerakis, and Dingzeyu Li. 2020. Makeltalk: speaker-aware talking-head animation. *ACM Transactions On Graphics* 39, 6 (2020), 1–15.
- [64] Hao Zhu, Huaibo Huang, Yi Li, Aihua Zheng, and Ran He. 2020. Arbitrary talking face generation via attentional audio-visual coherence learning. In *Proc. IJCAI*.
- [65] Hao Zhu, Man-Di Luo, Rui Wang, Ai-Hua Zheng, and Ran He. 2021. Deep audio-visual learning: A survey. *International Journal of Automation and Computing* 18 (2021), 351–376.