

FaceChain-ImagineID: Freely Crafting High-Fidelity Diverse Talking Faces from Disentangled Audio

Chao Xu^{1*} Yang Liu^{1*} Jiazheng Xing² Weida Wang² Mingze Sun²
 Jun Dan² Tianxin Huang³ Siyuan Li² Zhi-Qi Cheng⁴ Ying Tai⁵ Baigui Sun^{1†}

¹Alibaba Group ²FaceChain Community ³National University of Singapore

⁴Carnegie Mellon University ⁵Nanjing University

{xc264362, ly261666, baigui.sbg}@alibaba-inc.com




Figure 1. We propose an intuitive new paradigm, Listening and Imagining. *Listening* refers to the progressive disentanglement of identity, content, and emotion from the audio. *Imagining* involves generating diverse videos that align with the given decoupled audio cues while maintaining high consistency within each video. Notably, audio-unrelated attributes, e.g., beard, hairstyle, and pupil color, can be freely changed by specified prompts according to personal preferences (Sample 2). Other head pose and eye blink are copied from the reference videos. The audio is the only input to the model, while the left three images are just shown for facial geometry and semantic reference.

Abstract

In this paper, we abstract the process of people hearing speech, extracting meaningful cues, and creating various dynamically audio-consistent talking faces, termed *Listening and Imagining*, into the task of high-fidelity diverse talking faces generation from a single audio. Specifically, it involves two critical challenges: one is to effectively decouple identity, content, and emotion from entangled audio, and the other is to maintain intra-video diversity and inter-video consistency. To tackle the issues, we first dig out the intricate relationships among facial factors and simplify the decoupling process, tailoring a *Progressive Audio Disentanglement* for accurate facial geometry and semantics learning, where each stage incorporates a customized training module responsible for a specific factor. Secondly, to achieve visually diverse and audio-synchronized animation solely from input audio within a single model, we introduce the *Controllable Coherent Frame generation*, which involves the flexible integration of three trainable adapters with frozen *Latent Diffusion Models (LDMs)* to focus on maintaining facial geometry and semantics, as well as tex-

ture and temporal coherence between frames. In this way, we inherit high-quality diverse generation from LDMs while significantly improving their controllability at a low training cost. Extensive experiments demonstrate the flexibility and effectiveness of our method in handling this paradigm. The codes will be released at [FaceChain](#).

1. Introduction

Talking face generation [14, 51, 52, 55, 56, 63, 67] is a challenging task that aims to synthesize video based on provided audio and image. This technique finds wide application in various practical scenarios, especially virtual interaction. However, users encountered a dilemma during the process. They are concerned about privacy breaches when using real facial images, while the virtual avatars generated by off-the-shelf methods often fail to align well with their own voices. Thus, we envision a new paradigm if it is possible to liberate the specified source face and directly infer the synchronized virtual portrait that matches the real audio input.

Indeed, it is an intuitive process, as people often analyze a voice and then mentally visualize the corresponding video clip. To realize such Listening and Imagining paradigm, there are two critical issues. 1) How to disentangle face-related features *solely* from audio to ensure strong consistency between synthesized videos and given audio. We first explore the natural association between auditory and visual perception. It is true that the facial identity features are closely related to voice characteristics [2, 8, 40]. For example, a pronounced chin and prominent brow ridges usually accompany a deep voice, while women and children often have higher pitches. In addition, spoken content involves localized lip movement, and the emotion style reflects the global facial cues [11, 24]. Accurate decoupling of the above three factors, *i.e.*, identity, content, and emotion, has a significant impact on subsequent generation. However, existing researches [20, 36] mainly focus on the disentanglement of the latter two elements, while other studies [33, 59], although exploring identity, are geared towards generating static face image instead of dynamic animations. 2) How can we maintain diversity between different videos while ensuring consistency within each video by a *single* network. It is the consensus that human imagination is boundless. With the same audio, we can imagine numerous different talking videos, but within each clip, all frames share the uniform information. Benefiting from the progress of diffusion models, we have two potential methods available. One is the combination of text-to-image synthesis, *e.g.*, Latent Diffusion Models (LDMs) [42] and common driving methods like SadTalker [67] and DiffTalk [46]. The other is producing videos by borrowing the framework of text-to-video synthesis [12, 28, 38, 61]. However, the former involves two isolated models, while the latter is difficult to perfectly match the control conditions, resulting in noticeable differences between each frame. Additionally, neither of them fully utilizes the audio features.

In this paper, we are dedicated to solving these two problems. Firstly, due to the high degree of coupling in audio, it is challenging to directly employ cross reconstruction under the guidance of pseudo training pairs for disentanglement [20, 36]. To simplify it, we turn to the prior knowledge 3DMM for help and propose a Progressive Audio Disentanglement (PAD) to gradually decouple identity, content, and emotion. As shown in Fig. 1, we start by predicting the most independent and fundamental identity cues, including the explicit facial geometry and implicit semantics, *i.e.*, shape, gender, and age. Then the estimated shape is used as a condition to disentangle the localized content from the audio. We further supplement the remaining details of global emotion to generate comprehensive facial representations. In this way, we accurately estimate the face-related geometry and semantics from the audio. For the second concern, we propose a Controllable Coherent Frame (CCF) generation,

which offers several appealing designs. Firstly, inspired by the subject-driven models [43, 48], we combine the implicit audio cues with pred-defined prompt and map them to the CLIP domain [39], which influence the inferred faces both on semantics and emotions. Secondly, to avoid introducing extra offline components and ensure that the diffusion model has both invariance and variability, we freeze the LDMs and propose three trainable adapters to inject hierarchical conditions, which are equipped with a flexible mechanism that determines whether the outputs have new appearance or remain faithful to the neighboring frame. We further collaborate the CCF with autoregressive inference for diverse and temporally consistent video generation.

In summary, we present the following contributions.

- We propose a new paradigm, Listening and Imagining, for generating diverse and coherent talking faces based on a single audio, which is the first attempt in this field.
- We propose a novel PAD that gradually separates the identity, content, and emotion from the audio, providing a solid foundation for accurate and vivid face animations.
- We propose a novel frozen LDMs-based CCF with three trainable adapters to faithfully integrate the audio features and address the dilemma of intra-video diversity and inter-video consistency within a single model.

2. Related Work

2.1. Audio-driven Talking Face Generation

Talking face synthesis could be roughly divided into two folds. One of the research directions is GAN-based methods. Early efforts project the audio cues into latent feature space [6, 27, 35, 54, 55, 72, 73] and utilize a conditioned image generation framework to synthesize faces. To compensate for information loss in implicit codes, subsequent works have incorporated explicit structural information, such as landmark [46, 71] and 3DMM [41, 63, 65, 67], to more accurately reflect the audio features on the face. Similarly, 3DMM is used as an intermediate representation in our method. We further leverage it to fully decouple the audio signal. Motivated by progress of the diffusion model in image synthesis tasks, recent approaches [10, 46, 51, 64] focus on another research line. They carefully design conditional control modules and train them along with UNet to achieve high-fidelity faces. However, trainable diffusion models fail to inherit many appealing properties, such as diverse generation. In this work, we explore the feasibility of using pretrained models to maintain diversity.

2.2. Audio to Face Generation

Human voices carry a large amount of personal information, including speaker identity [8, 40], age [15, 49], gender [25] and emotion style [58, 69]. Based on previous studies, it is possible to directly predict entire face from corresponding

audio. Common methods [1, 4, 33, 59] adopt GANs to generate face images from voice input. However, the above reconstruction-based attempts are not reasonable because audio lacks specific visual features like hairstyles and backgrounds, while these attributes can vary significantly within the same person. Consequently, CMP [60] only models the correlation between facial geometry and voice, but the results lack authenticity. We propose an ideal solution that involves inferring facial structure information from audio and then combining it with a conditional generative model to achieve controllable and diverse generation.

2.3. Diffusion Models

Diffusion models [18, 32] are popular for generating realistic and diverse samples. In practice, DDIM [50] converts the sampling process to a faster non-Markovian process, while LDMs [42] perform diffusion in a compressed latent space to reduce memory and runtime. Recently, text-to-image generations [13, 19, 26, 44, 45] have gained significant attention. For personalized control, DreamBooth [43] proposes a subject-driven generation by fine-tuning diffusion models. To reduce training costs, T2I-Adapter [31] only trains the extra encoder to influence the denoising process. Similarly, our work incorporates the above two characteristics. Another challenging topic is text-to-video generation [3, 12, 61, 62, 68]. However, they suffer from inconsistency across video frames both temporally and semantically. Additionally, they struggle to generate long videos, impacting the production of high-fidelity talking faces. Therefore, we utilize an autoregressive inference strategy to ensure smooth and consistent transitions between frames, even for long-duration videos.

3. Method

Listening and imagining are instinctive behaviours for humans to perceive and comprehend the world. When people hear a voice, they first analyze the characteristics of the speaker, *i.e.*, who is speaking (Identity), what is being said (Content), and how it feels (Emotion), and then imagine diverse and temporally consistent dynamic scenes that match the inferred audio features. To simulate such process, we follow a two-stage phase. Firstly, a Progressive Audio Disentanglement (PAD) is designed to gradually separate the identity, content, and emotion from the coupled audio in Sec. 3.1. Secondly, we design Controllable Coherent Frame (CCF) generation (Sec. 3.2), which inherits the diverse generation capabilities from Latent Diffusion Models (LDMs) and develops several techniques to enhance controllability: Texture Inversion Adapter (TIA), Spatial Conditional Adapter (SCA), and Mask-guided Blending Adapter (MBA) jointly ensure the generated result meets specific conditions. Autoregressive Inference is utilized to maintain temporal consistency throughout the video.

3.1. Progressive Audio Disentanglement

To extract independent cues from the entangled audio signals, previous work [20] builds aligned pseudo pairs and adopts the cross-reconstruction manner for content and emotion disentanglement. However, due to the inevitable introduction of errors when constructing pseudo pairs, they become unreliable when more elements are involved. To this end, we only use a single audio for training and adopt 3DMM [9] to achieve accurate disentangled 3D facial prior, *e.g.*, shape $\alpha \in \mathbb{R}^{80}$, expression sequences $\beta \in \mathbb{R}^{L \times 64}$, *etc.*, which excludes the face-unrelated factors, such as hairstyle and makeup, and serve as the ground truth for supervision during training. Note that a video shares the same shape while the expression corresponds to each frame. In particular, we propose *Progressive Audio Disentanglement* to simplify the decoupling process by extracting the fundamental identity information, then separating the localized mouth movements, and finally obtaining the global expression, outlined in Fig. 2. For more discussions about the disentanglement order (identity \rightarrow content \rightarrow emotion), the superiority of PAD compared to other methods, and architecture details, please refer to the supplementary materials.

Identity. Motivated by the recent studies that voice is articulatory related to skull shape, as well as gender and age [16, 29], the identity hinted in audio could refer to the facial geometry and semantic embedding. In practice, the identity encoder Φ_E^{id} is the Transformer-based architecture with an extended learnable token $\tilde{\alpha} \in \mathbb{R}^{512}$ of face shape, mapping the MFCC sequence $\mathbf{A} \in \mathbb{R}^{L \times 1280}$ to both the estimated shape $\hat{\alpha} \in \mathbb{R}^{80}$ and semantics $\theta_{id} \in \mathbb{R}^{512}$:

$$\tilde{\alpha}, \bar{\theta}_{id} = \Phi_E^{id}([\tilde{\alpha}, \text{MLPs}(\mathbf{A})] + \text{PE}), \quad (1)$$

$$\hat{\alpha} = \text{MLPs}(\tilde{\alpha}), \theta_{id} = \text{Avg}(\bar{\theta}_{id}), \quad (2)$$

where PE is the positional embedding, $\tilde{\alpha} \in \mathbb{R}^{512}$, $\bar{\theta}_{id} \in \mathbb{R}^{L \times 512}$. During training, we calculate Reconstruction Loss between predicted $\hat{\alpha}$ and the ground truth α : $\mathcal{L}_{id} = \|\alpha - \hat{\alpha}\|_2$. Moreover, we apply Contrastive Loss on semantic embeddings by constructing a positive pair $(\theta_{id}^p, \theta_{id})$ with the same identity and a negative pair $(\theta_{id}^n, \theta_{id})$ with the different. Then InfoNCE loss [34] with cosine similarity \mathcal{S} is enhanced between two pairs:

$$\mathcal{L}_{con} = -\log \left[\frac{\exp(\mathcal{S}(\theta_{id}^p, \theta_{id}))}{\exp(\mathcal{S}(\theta_{id}^p, \theta_{id})) + \exp(\mathcal{S}(\theta_{id}^n, \theta_{id}))} \right]. \quad (3)$$

Content. The 3DMM expression coefficient involves local lip and global facial movements, while the spoken content is mainly related to the former. Inspired by the work [55], we employ the lip reading expert [47] into the training phase and produce the only mouth-related expression coefficients $\mathbf{l} \in \mathbb{R}^{L \times 64}$ for content disentanglement. As shown in Fig. 2, we construct a Transformer-based encoder-decoder architecture, where the linguistic features $\bar{\mathbf{l}} \in \mathbb{R}^{L \times 64}$ embedded

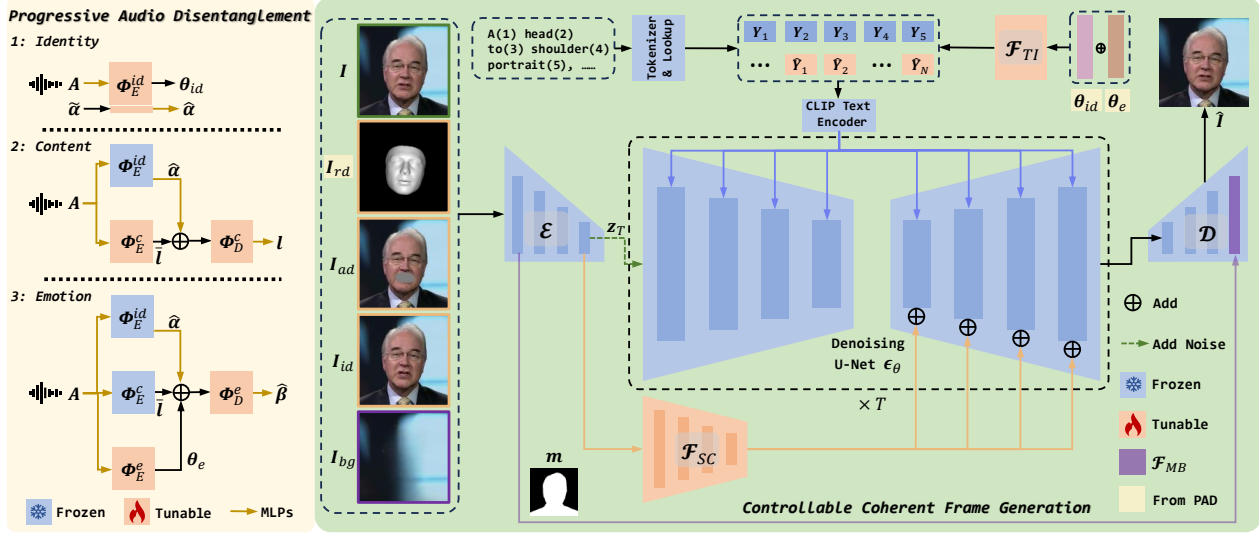


Figure 2. **Overview of the proposed method.** Our approach involves a two-stage framework that corresponds to the Listening and Imagining. For listening, the Progressive Audio Disentanglement gradually separates the identity, content, and emotion from the entangled audio. For Imagining, Controllable Coherent Frame generation receives the facial semantics (θ_{id} and θ_e inferred from PAD) and geometry (3D mesh I_{rd} rendered from $\hat{\alpha}$, $\hat{\beta}$ and other coefficients extracted from I) to synthesize the diverse audio-synchronized faces, while the I_{ad} , I_{id} , and I_{bg} are further introduced to achieve highly controllable generation with complete visual and temporal consistency. Please refer to Alg. 1 for more details. In this way, we achieve diverse and high-fidelity face animation solely from audio.

by the content encoder Φ_E^c and the identity template output by the frozen identity encoder are combined together and sent into the content decoder Φ_D^c to obtain l :

$$\bar{l} = \Phi_E^c(\text{MLPs}(A) + \text{PE}), \quad (4)$$

$$l = \text{MLPs}(\Phi_D^c(\bar{l} + \text{MLPs}(\hat{\alpha}))). \quad (5)$$

We train this stage by using two losses: Regularization Loss calculates the distance between β and l with a relative small weight to smooth training phase: $\mathcal{L}_{reg} = \|l - \beta\|_2$. Besides, we project the coefficients to the image domain with textures, and utilize lip-reading expert to predict the text \hat{X} . Assuming that the original text content is X , we compute the Lip-reading Loss via cross-entropy: $\mathcal{L}_{lip} = -X \log P(\hat{X}|V)$, where V is the ground truth video.

Emotion. In this stage, we directly utilize β as constraint to decouple the remaining emotion styles. As shown in Fig. 2, based on the second stage, an additional emotion encoder Φ_E^e is introduced to generate the pooled emotion embeddings $\theta_e \in \mathbb{R}^{512}$, which is combined with identity $\hat{\alpha}$ and linguistic features \bar{l} for the expression coefficient prediction by the emotion decoder Φ_D^e :

$$\theta_e = \text{Avg}(\Phi_E^e(\text{MLPs}(A) + \text{PE})), \quad (6)$$

$$\hat{\beta} = \text{MLPs}(\Phi_D^e(\theta_e + \bar{l} + \text{MLPs}(\hat{\alpha}))). \quad (7)$$

Similar to Eq. 3, we adopt Contrastive Loss to extract more discriminative emotion features, details of which are omitted here. Additionally, the regular Reconstruction Loss is used to supervise generated $\hat{\beta}$.

3.2. Controllable Coherent Frame Generation

Existing methods typically specify the input source face along with audio cues for video generation, while our objective is to achieve visually diverse and audio-consistent animation directly from input audio. Recent LDMs are naturally suited to diverse generation, and their conditional text prompts offer a pathway for freely editing attributes that cannot be deduced from the audio, as shown in Fig. 1. In exchange, they are relatively weak in controllability, especially in video generation [38, 61]. Thus, to tailor it to our task without introducing extra offline models (e.g., LDMs + DiffTalk[46]), we must tackle two challenges: ensuring that the synthesized video content aligns with the given conditions, and achieving smooth temporal transitions across frames. Our targeted designs are depicted in Fig. 2.

Textual Inversion Adapter. The critical issue in injecting the identity θ_{id} and emotion θ_e inferred from PAD to the frozen LDMs lies in aligning them with the CLIP domain. Inspired by the inversion technique [43], we propose a *Textual Inversion Adapter* \mathcal{F}_{TI} , which maps the input vector-based conditions into a set of token embeddings to sufficiently represent these conditions in the CLIP token embedding space. In detail, we first predefine the prompt for basic high-quality face generation. It is tokenized and mapped into the token embedding space by employing a CLIP embedding lookup module, obtaining $\{Y_1, \dots, Y_M\}$. Then, the adapter encodes θ_{id} and θ_e into the aligned pseudo-word token embeddings $\{\hat{Y}_1, \dots, \hat{Y}_N\}$. We concatenate these two embeddings and feed them to the CLIP text en-

Algorithm 1 Autoregressive Inference

Input: z_T : random noise
 $\{\mathbf{I}_{rd}^i\}_{i=1}^H$: 3D mesh sequences \triangleright Inferred from audio
 \mathbf{Y} : token embeddings \triangleright Inferred from audio

$\triangleright i = 1$, Diverse mode
 $\hat{\mathbf{I}}_1 = \mathcal{D}(\text{CCF}(z_T, \mathcal{E}(\mathbf{I}_{rd}^1), \mathbf{Y}))$

$\triangleright i = 2 \dots H$, Coherent mode
for $i = 2 \dots H$ **do**
 $\mathbf{I}_{id}^i = \hat{\mathbf{I}}_1, \mathbf{I}_{bg}^i = \mathcal{G}_{IA}(\hat{\mathbf{I}}_1)$
 $\mathbf{I}_{ad}^i = \mathcal{M}(\hat{\mathbf{I}}_{i-1}), \mathbf{m}_i = \mathcal{G}_{MOD}(\hat{\mathbf{I}}_{i-1})$
 $\hat{\mathbf{I}}^i = \mathcal{D}'(\text{CCF}(z_T, \mathcal{E}([\mathbf{I}_{rd}^i, \mathbf{I}_{id}^i, \mathbf{I}_{ad}^i]), \mathbf{Y}), \mathcal{E}(\mathbf{I}_{bg}^i), \mathbf{m}_i)$
end for
Output: $\hat{\mathbf{V}}$: generated video

coder, whose output \mathbf{Y} is applied on the cross-attention layer of LDMs to guide the denoising process. When training this adapter, we freeze all of the other parameters:

$$\mathbb{E}_{z_0, \varepsilon \sim N(0, \mathbf{I}), t, \mathbf{Y}} \|\varepsilon - \varepsilon_\theta(z_t, t, \mathbf{Y})\|_2^2. \quad (8)$$

Spatial Conditional Adapter. Inspired by T2I-Adapter, we devise *Spatial Conditional Adapter* \mathcal{F}_{SC} to further fuse the explicit conditions. As shown in Fig. 2, 3D face mesh \mathbf{I}_{rd} contains rich audio-synchronized facial geometry, *i.e.*, face shape, lip movement, and expression style. Besides, we sample a random frame \mathbf{I}_{id} of the same identity to provide the face appearance and background. These two conditions are enough for common methods [67, 73] to animate the identity face to desired pose and expression. However, it is difficult for frozen LDMs to learn complex spatial transformation. Thus, we further incorporate the adjacent frame \mathbf{I}_{ad} as an additional reference, and mask its mouth area \mathcal{M} produced by 3D mesh to avoid the networks from learning shortcuts by directly copying from it. This condition makes the deformation learning much easier and also provides motion cues. We only train this adapter in this stage:

$$\mathbb{E}_{z_0, \varepsilon \sim N(0, \mathbf{I}), t, \mathbf{F}_{sc}} \|\varepsilon - \varepsilon_\theta(z_t, t, \mathbf{Y}, \mathbf{F}_{sc})\|_2^2, \quad (9)$$

where $\mathbf{F}_{sc} = \mathcal{F}_{SC}(\mathcal{E}(\mathbf{I}_{rd}) + \mathcal{E}(\mathbf{I}_{id}) + \mathcal{E}(\mathbf{I}_{ad}))$ when $p < 0.8$, else $\mathcal{F}_{SC}(\mathcal{E}(\mathbf{I}_{rd}))$. \mathcal{E} is the VAE encoder. p is a random number generator within the range from 0 to 1.

Masked-guided Blending Adapter. Despite spatial pixel-level conditions have been given, we observe significant distortion artifacts in the background, resulting in a low fidelity of the synthesized video. Therefore, we introduce a *Masked-guided Blending Adapter* \mathcal{F}_{MB} into VAE Decoder \mathcal{D} , forming \mathcal{D}' , which utilizes a mask to directly copy the masked region of the k -th decoder feature \mathbf{F}_{vd}^k while combining the unmasked region of the k -th background feature \mathbf{F}_{ve}^k from the VAE encoder. The dilated mask \mathbf{m} and background \mathbf{I}_{bg} are produced by MODNet [23] \mathcal{G}_{MOD} and




Figure 3. Qualitative results of audio-to-face on MEAD. Icons of the same color indicate samples from the same audio. Ours-I, -C, and -E mean the stage of identity, content, emotion decoupling.

Inpaint-Anything [66] \mathcal{G}_{IA} respectively. The blending process is formulated as:

$$\hat{\mathbf{F}}_{vd}^k = \mathbf{F}_{vd}^k \otimes \mathbf{m} + \text{Conv}(\mathbf{F}_{ve}^k) \otimes \tilde{\mathbf{m}}, \quad (10)$$

where \otimes is element-wise multiplication, $\tilde{\mathbf{m}} = 1 - \mathbf{m}$. Experimentally, we inset this module into the decoder layer with the resolution of 512. The training only involves a frozen autoencoder and a trainable convolution layer, under the supervision of reconstruction and VGG losses. With this design, the synthesized video maintains a consistent background, and the fusion edge is seamlessly harmonious. Note that the above three adapters are trained sequentially.

Autoregressive Inference. The designed CCF has diverse and controllable image generation capabilities. We incorporate an *Autoregressive Inference* to further generate diverse and temporally coherent videos. As shown in Alg. 1, unlike DiffTalk [46] starting from a given frame, for the first frame, we switch CCF to diverse mode, *i.e.*, only receive audio-enhanced CLIP embedding \mathbf{Y} and encoded 3D face mesh features $\mathcal{E}(\mathbf{I}_{rd}^1)$, generating faces $\hat{\mathbf{I}}_1$ with varying appearance while maintaining the consistent geometry with the audio. For the subsequent frames, we further apply three adapters to achieve controllable generation. Specifically, \mathbf{I}_{id}^i and \mathbf{I}_{bg}^i are $\hat{\mathbf{I}}_1$ and $\mathcal{G}_{IA}(\hat{\mathbf{I}}_1)$ all the time, \mathbf{I}_{rd}^i is the mesh at time stamp i , and \mathbf{I}_{ad}^i is the masked former frame $\hat{\mathbf{I}}_{i-1}$, $\mathbf{m}_i = \mathcal{G}_{MOD}(\hat{\mathbf{I}}_{i-1})$ since the difference between adjacent frames is small.



Figure 4. Visual comparison with recent SOTA methods. Images are from officially released codes for fair comparisons. The first sample selected from HDTF, second from MEAD. For the third, based on the audio of the first row, our method generate a unseen face for all competitors, which using this as the source face to produce talking faces. The first row provides the ground truth for facial expression.

4. Experiments

4.1. Experimental Setup

Datasets. We adopt three talking face dataset MEAD [57], VoxCeleb2 [7], and HDTF [70] in our experiments.

Metrics. We adopt PSNR, SSIM, and FID [17] to measure the image quality of synthesized videos. The distance between the landmarks of the mouth (LMD) and confidence score (Sync) proposed in SyncNet [5] is used to measure the audio-visual synchronization. Besides, we compute emotion accuracy (EACC) to measure the generated emotions. Furthermore, we compute the cosine similarity of geometry (Shape) and semantics (Age, Gender, Emotion), and overall quality (FID) between the ground truth and generated faces inferred from audio by using several off-the-shelf tools, D3DFR [9], FairFace [22], and FAN [30].

Implementation Details. For PAD training, we adopt VoxCeleb2 for identity disentanglement and MEAD training set for content and emotion. The audios are pre-processed to 16kHz and converted to mel-spectrogram with 80 Mel filterbank, as Wav2Lip. The length of the sampled video clip is $L = 32$. For CCF training, the input videos are cropped and resized to 512×512 . We adopt MEAD and HDTF as training sets and optimize TIA, SCA, and MBA sequentially. It takes about 2.5 days totally when trained on 8 V100 GPUs.

$N = 8$ and M is the length of the prompts in SCA. During quantitative assessment, we sample 12 identities from MEAD and 10 from HDTF, which serve as the test set. Denoising step T is set as 50 for both training and inference.

4.2. Comparison with State-of-the-Art Methods.

Qualitative Results. In this section, we first perform audio-to-face method comparisons with CMP [60]. As shown in Fig. 3, we display the generated 3D meshes and their corresponding faces of our method in rows 4, 5, and CMP in rows 6, 7. The real faces in the first row provide geometry and semantics references. It can be seen that CMP fails to generate desired shapes and the synthesized faces have blurred details with obvious artifacts. By contrast, benefiting from progressive disentanglement and LDMs, our method produces more accurate geometry, including shape, lip movement, emotion styles, and high-quality realistic textures. Besides, in each sample, the first two columns correspond to the same audio, while the third represents a different audio of the same person. Across different audio clips, our method produces more consistent results than CMP. We further compare our method with talking face methods including Wav2Lip, PC-AVS, EAMM, and SadTalker, which are reproduced from their officially released codes. As shown in Fig. 4, in terms of emotion accuracy, audio-visual syn-




Figure 5. Illustration of disentangled controllability. (a) is under the diverse mode and (b) is under the coherent mode.

Method	Shape \uparrow	Age \uparrow	Gender \uparrow	Emotion \uparrow	FID \downarrow
CMP [60]	0.60	0.41	0.87	0.23	228.05
Ours	0.75	0.58	0.93	0.76	25.78

Table 1. Quantitative comparison with CMP. The value is the average of all samples on two test sets (same as the Tab. 2).

Method	EACC \downarrow	LMD \downarrow	Sync \uparrow	FID \downarrow	PSNR \uparrow	SSIM \uparrow
Wav2Lip [37]	0.129	2.96	5.93	40.79	29.05	0.59
PC-AVS [73]	0.102	<u>2.71</u>	5.75	45.84	29.28	0.59
EAMM [21]	<u>0.097</u>	2.75	5.46	50.55	29.14	0.61
SadTalker [67]	0.121	2.80	5.84	31.62	29.55	<u>0.64</u>
Ours-C	0.113	2.82	5.86	28.09	<u>29.59</u>	0.60
Ours-E	0.090	2.67	<u>5.91</u>	<u>28.37</u>	29.66	0.67

Table 2. Quantitative comparison with state-of-the-art talking face methods on two test sets.

chronization, and image quality, our method outperforms these SOTA methods. Notably, for the third sample, other competitors adhere the two stage paradigm, *i.e.*, Audio-to-Image and then Image-to-Video, while our method autoregressively infers video from the given audio.

Quantitative Results. As shown in Tab. 1, our method excels over CMP across all metrics, especially in FID, which is consistent with the qualitative results in Fig. 3. Moreover, we observe that gender is relatively easy to predict, while others are more challenging. Yet, our method still gains significant improvements. We further present the quantitative comparison in Tab. 2 of talking face generation, which reveals our method can achieve the best performance in most metrics. Wav2Lip surpasses all the competitors on Sync due to the supervision of the synchronization scoring model. We further report the user study in supplementary materials.

4.3. Further Analysis

Disentangled Controllability. To demonstrate the facial factor disentanglement of our method, we first sample two distinctly different faces and exchange one factor at a time in diverse mode. Although our method independently controls the desired factor, *i.e.*, content (lip movement), semantics (gender and age), and emotion styles in columns




Figure 6. Illustration of the image retrieval using the identity semantic features. The faces (column 1) only shown for reference.




Figure 7. Clusters of the disentangled emotion embeddings. Different emotions are clustered in different regions.

3, 4, 5 of Fig. 5(a), the results inherit the weaknesses of the LDM in preserving visual consistency. As the red rectangles marked in row 1, despite only editing the mouth shape, they exhibit noticeable difference in appearance, which indicates that simply generating frames and splicing them together cannot create visually and temporally consistent videos. Furthermore, we conduct a more rigorous analysis in coherent mode. The emotion control (cols. 1-4) and content editing (cols. 4-5) are illustrated in Fig. 5(b), verifying the effectiveness of the disentanglement.

Interpretability of Disentanglement. We further perform qualitative experiments to demonstrate the disentanglement of learned identity and emotion representation. Firstly, we randomly sample four audios of different ages and genders from the CelebV-HQ [74] dataset and select their most consistent pairs according to the cosine similarity between two semantic features. We display the corresponding images for visual comparisons in Fig. 6, where the retrieved audios have similar ages and genders to those of the query. We also attach the generated face in column 2, which closely resembles the conditions in column 1, illustrating the effectiveness of our method. Secondly, to evaluate the disentanglement of the emotion, we use t-SNE [53] to visualize the




Figure 8. Ablation study on the number of pseudo-word tokens in TIA. We do not apply SCA and MBA in this stage.



Figure 9. Ablation study on the design of the conditions in SCA.

emotion latent space in Fig. 7. It is obvious that the different emotions are clustered into separate groups. Based on these analyses, we conclude that the PAD contributes to the decoupling entangled facial cues from the audio.

4.4. Ablation Study and Efficiency Evaluation

Impact of Disentangled Module in PAD. To verify the superiority of PAD, we show the generated meshes under different stages in Fig. 3 rows 2-4. The first stage handles identity disentanglement, and the corresponding meshes accurately reflect desired face shapes and outlines. The following content disentanglement ensures that the mouth movements synchronize with the spoken phonemes. The final emotion disentanglement enhances the facial global cues, leading to highly consistent meshes with the reference faces, *e.g.*, more accurate lips with fearful emotion (the right sample). Note that the left sample with neutral shows no noticeable changes. Besides, the last two rows of Tab. 2 can also prove the effectiveness of PAD.

Impact of Pseudo-word Tokens in TIA. In Fig. 8, we show the results when varying the number of pseudo-word tokens in TIA. We observe that too few tokens can not synthesize semantically consistent results, while too many do not lead to further visual improvement. To balance the computational load and performance, we set N as 8 experimentally. Besides, the sixth column indicates that when directly inputting audio features into TIA, the output suffers degradation due to insufficient decoupling. We further attach the quantitative results in Fig. 8. Since TIA does not involve face geometry, shape metric is not displayed. These results are based on the 100 faces sampled from CelebV-HQ.

Impact of Conditions in SCA. The conditions in SCA are the critical design in this work for generating control-




Figure 10. Ablation study on the design of the blending structure in MBA. This case is sampled from Fig. 4.

lable and coherent frames. As shown in Fig. 9 row 1, the model with TIA can generate semantically consistent faces but fails to preserve the geometrical information (column 2). Introducing the SCA module with I_{rd} input tackles the spatial misalignment (column 3). The adjacent frame I_{ad} is further incorporated trying to maintain the appearance like the reference face (column 4), but it introduces obvious artifacts, especially in the mouth area. When taking I_{rd} and I_{id} as input, as previously mentioned, the frozen LDMs encounter difficulties in learning complex transformations, resulting in both the appearance and geometry are not well aligned (column 5). In this work, we employ all three inputs together for SCA (column 6), where I_{rd} is responsible for the facial structure, I_{ad} for the appearance, while I_{id} further refines the details and enhances the image quality.

Impact of Blending Structure in MBA. In Fig. 10, the result in the second column is produced by the method without MBA, which fails to preserve the background details. To solve this challenge, we design several variants to explore the blending structure. Without mask guidance, the synthesized background still differs from the ground truth (column 3). For this reason, we introduce the mask into all VAE decoder layers (column 4). Due to the edge aliasing in the low-resolution mask, although it solves the background issue, obvious artifacts appear around the blending edges. When applied at 512 resolution (column 6), the generated face achieves the best visual performance, achieving consistent background and seamless edge blending.

Impact of Training Strategy in CCF. We explore the effectiveness of sequential training in CCF. As shown in Fig. 9, when jointly training TIA and SCA, the former does not work well as facial semantics can take shortcuts, *i.e.*, learning from I_{id} and I_{ad} instead of token embeddings Y .

Efficiency Evaluation. We follow the DiffTalk to open the UNet parameters during SCA training, the number of trainable parameters increases from 282.9M to 1.15G. Furthermore, tuning the UNet does not bring significant benefits in Fig. 9. Thus, our method is effective and user-friendly.

5. Conclusions

We propose a novel framework to simulate the process of Listening and Imagining. Firstly, PAD simplifies the decoupling process and improves the synchronization of predicted representations with audio. Secondly, CCF further generates realistic and diverse talking faces. Extensive experiments demonstrate the effectiveness of our approach.

References

- [1] Yeqi Bai, Tao Ma, Lipo Wang, and Zhenjie Zhang. Speech fusion to face: Bridging the gap between human’s vocal characteristics and facial imaging. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 2042–2050, 2022. **3**
- [2] Ray Bull, Harriet Rathborn, and Brian R Clifford. The voice-recognition accuracy of blind listeners. *Perception*, 12(2): 223–226, 1983. **2**
- [3] Weifeng Chen, Jie Wu, Pan Xie, Hefeng Wu, Jiashi Li, Xin Xia, Xuefeng Xiao, and Liang Lin. Control-a-video: Controllable text-to-video generation with diffusion models. *arXiv preprint arXiv:2305.13840*, 2023. **3**
- [4] Hyeong-Seok Choi, Changdae Park, and Kyogu Lee. From inference to generation: End-to-end fully self-supervised generation of human face from speech. *arXiv preprint arXiv:2004.05830*, 2020. **3**
- [5] Joon Son Chung and Andrew Zisserman. Out of time: automated lip sync in the wild. In *Computer Vision—ACCV 2016 Workshops: ACCV 2016 International Workshops, Taipei, Taiwan, November 20–24, 2016, Revised Selected Papers, Part II 13*, pages 251–263. Springer, 2017. **6**
- [6] Joon Son Chung, Amir Jamaludin, and Andrew Zisserman. You said that? *arXiv preprint arXiv:1705.02966*, 2017. **2**
- [7] Joon Son Chung, Arsha Nagrani, and Andrew Zisserman. Voxceleb2: Deep speaker recognition. *arXiv preprint arXiv:1806.05622*, 2018. **6**
- [8] Angus Deaton. Understanding the mechanisms of economic development. *Journal of Economic Perspectives*, 24(3):3–16, 2010. **2**
- [9] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019. **3, 6**
- [10] Chenpeng Du, Qi Chen, Tianyu He, Xu Tan, Xie Chen, Kai Yu, Sheng Zhao, and Jiang Bian. Dae-talker: High fidelity speech-driven talking face generation with diffusion auto-encoder. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 4281–4289, 2023. **2**
- [11] Paul Ekman and Wallace V Friesen. Facial action coding system. *Environmental Psychology & Nonverbal Behavior*, 1978. **2**
- [12] Patrick Esser, Johnathan Chiu, Parmida Atighehchian, Jonathan Granskog, and Anastasis Germanidis. Structure and content-guided video synthesis with diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7346–7356, 2023. **2, 3**
- [13] Aditya Ramesh et al. Hierarchical text-conditional image generation with clip latents, 2022. **3**
- [14] Yue Gao, Yuan Zhou, Jinglu Wang, Xiao Li, Xiang Ming, and Yan Lu. High-fidelity and freely controllable talking head video generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5609–5619, 2023. **1**
- [15] Joanna Grzybowska and Stanislaw Kacprzak. Speaker age classification and regression using i-vectors. In *INTER-SPEECH*, pages 1402–1406, 2016. **2**
- [16] William J Hardcastle, John Laver, and Fiona E Gibbon. *The handbook of phonetic sciences*. John Wiley & Sons, 2012. **3**
- [17] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. **6**
- [18] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. **3**
- [19] Lianghua Huang, Di Chen, Yu Liu, Yujun Shen, Deli Zhao, and Jingren Zhou. Composer: Creative and controllable image synthesis with composable conditions. *arXiv preprint arXiv:2302.09778*, 2023. **3**
- [20] Xinya Ji, Hang Zhou, Kaisiyuan Wang, Wayne Wu, Chen Change Loy, Xun Cao, and Feng Xu. Audio-driven emotional video portraits. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14080–14089, 2021. **2, 3**
- [21] Xinya Ji, Hang Zhou, Kaisiyuan Wang, Qianyi Wu, Wayne Wu, Feng Xu, and Xun Cao. Eamm: One-shot emotional talking face via audio-based emotion-aware motion model. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10, 2022. **7**
- [22] Kimmo Karkkainen and Jungseock Joo. Fairface: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1548–1558, 2021. **6**
- [23] Zhanghan Ke, Jiayu Sun, Kaican Li, Qiong Yan, and Rynson WH Lau. Modnet: Real-time trimap-free portrait matting via objective decomposition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1140–1147, 2022. **5**
- [24] Dacher Keltner, Disa Sauter, Jessica Tracy, and Alan Cowen. Emotional expression: Advances in basic emotion theory. *Journal of nonverbal behavior*, 43:133–160, 2019. **2**
- [25] Sheng Li, Raj Dabre, Xugang Lu, Peng Shen, Tatsuya Kawahara, and Hisashi Kawai. Improving transformer-based speech recognition systems with compressed structure and speech attributes augmentation. In *Interspeech*, pages 4400–4404, 2019. **2**
- [26] Yuheng Li, Haotian Liu, Qingyang Wu, Fangzhou Mu, Jianwei Yang, Jianfeng Gao, Chunyuan Li, and Yong Jae Lee. Gligen: Open-set grounded text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22511–22521, 2023. **3**
- [27] Borong Liang, Yan Pan, Zhizhi Guo, Hang Zhou, Zhibin Hong, Xiaoguang Han, Junyu Han, Jingtuo Liu, Errui Ding, and Jingdong Wang. Expressive talking head generation with granular audio-visual control. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3387–3396, 2022. **2**
- [28] Yue Ma, Yingqing He, Xiaodong Cun, Xintao Wang, Ying Shan, Xiu Li, and Qifeng Chen. Follow your pose:

- Pose-guided text-to-video generation using pose-free videos. *arXiv preprint arXiv:2304.01186*, 2023. 2
- [29] Dieter Maurer and Theodor Landis. Role of bone conduction in the self-perception of speech. *Folia phoniatrica*, 42(5): 226–229, 1990. 3
- [30] Debin Meng, Xiaojiang Peng, Kai Wang, and Yu Qiao. frame attention networks for facial expression recognition in videos. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 3866–3870. IEEE, 2019. 6
- [31] Chong Mou, Xintao Wang, Liangbin Xie, Jian Zhang, Zhong-gang Qi, Ying Shan, and Xiaohu Qie. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. *arXiv preprint arXiv:2302.08453*, 2023. 3
- [32] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021. 3
- [33] Tae-Hyun Oh, Tali Dekel, Changil Kim, Inbar Mosseri, William T Freeman, Michael Rubinstein, and Wojciech Matusik. Speech2face: Learning the face behind a voice. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7539–7548, 2019. 2, 3
- [34] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 3
- [35] Youxin Pang, Yong Zhang, Weize Quan, Yanbo Fan, Xiaodong Cun, Ying Shan, and Dong-ming Yan. Dpe: Disentanglement of pose and expression for general video portrait editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 427–436, 2023. 2
- [36] Ziqiao Peng, Haoyu Wu, Zhenbo Song, Hao Xu, Xiangyu Zhu, Jun He, Hongyan Liu, and Zhaoxin Fan. Emotalk: Speech-driven emotional disentanglement for 3d face animation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20687–20697, 2023. 2
- [37] KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Namboodiri, and CV Jawahar. A lip sync expert is all you need for speech to lip generation in the wild. In *Proceedings of the 28th ACM international conference on multimedia*, pages 484–492, 2020. 7
- [38] Chenyang Qi, Xiaodong Cun, Yong Zhang, Chenyang Lei, Xintao Wang, Ying Shan, and Qifeng Chen. Fatezero: Fusing attentions for zero-shot text-based video editing. *arXiv preprint arXiv:2303.09535*, 2023. 2, 4
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2
- [40] Mirco Ravanelli and Yoshua Bengio. Speaker recognition from raw waveform with sinenet. In *2018 IEEE spoken language technology workshop (SLT)*, pages 1021–1028. IEEE, 2018. 2
- [41] Yurui Ren, Ge Li, Yuanqi Chen, Thomas H Li, and Shan Liu. Pirenderer: Controllable portrait image generation via semantic neural rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13759–13768, 2021. 2
- [42] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 2, 3
- [43] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22500–22510, 2023. 2, 3, 4
- [44] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Wei Wei, Tingbo Hou, Yael Pritch, Neal Wadhwa, Michael Rubinstein, and Kfir Aberman. Hyperdreambooth: Hypernetworks for fast personalization of text-to-image models. *arXiv preprint arXiv:2307.06949*, 2023. 3
- [45] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022. 3
- [46] Shuai Shen, Wenliang Zhao, Zibin Meng, Wanhua Li, Zheng Zhu, Jie Zhou, and Jiwen Lu. Diftalk: Crafting diffusion models for generalized talking head synthesis. *arXiv preprint arXiv:2301.03786*, 2023. 2, 4, 5
- [47] Bowen Shi, Wei-Ning Hsu, Kushal Lakhotia, and Abdelrahman Mohamed. Learning audio-visual speech representation by masked multimodal cluster prediction. *arXiv preprint arXiv:2201.02184*, 2022. 3
- [48] Jing Shi, Wei Xiong, Zhe Lin, and Hyun Joon Jung. Instantbooth: Personalized text-to-image generation without test-time finetuning. *arXiv preprint arXiv:2304.03411*, 2023. 2
- [49] Rita Singh, Joseph Keshet, Deniz Gencaga, and Bhiksha Raj. The relationship of voice onset time and voice offset time to physical age. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5390–5394. IEEE, 2016. 2
- [50] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 3
- [51] Michał Stypułkowski, Konstantinos Vougioukas, Sen He, Maciej Zieba, Stavros Petridis, and Maja Pantic. Diffused heads: Diffusion models beat gans on talking-face generation. *arXiv preprint arXiv:2301.03396*, 2023. 1, 2
- [52] Shuai Tan, Bin Ji, and Ye Pan. Emmn: Emotional motion memory network for audio-driven emotional talking face generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22146–22156, 2023. 1
- [53] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9 (11), 2008. 7
- [54] Duomin Wang, Yu Deng, Zixin Yin, Heung-Yeung Shum, and Baoyuan Wang. Progressive disentangled representation

- learning for fine-grained controllable talking head synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17979–17989, 2023. [2](#)
- [55] Jiadong Wang, Xinyuan Qian, Malu Zhang, Robby T Tan, and Haizhou Li. Seeing what you said: Talking face generation guided by a lip reading expert. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14653–14662, 2023. [1](#), [2](#), [3](#)
- [56] Jiayu Wang, Kang Zhao, Shiwei Zhang, Yingya Zhang, Yujun Shen, Deli Zhao, and Jingren Zhou. Lipformer: High-fidelity and generalizable talking face generation with a pre-learned facial codebook. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13844–13853, 2023. [1](#)
- [57] Kaisiyuan Wang, Qianyi Wu, Linsen Song, Zhuoqian Yang, Wayne Wu, Chen Qian, Ran He, Yu Qiao, and Chen Change Loy. Mead: A large-scale audio-visual dataset for emotional talking-face generation. In *European Conference on Computer Vision*, pages 700–717. Springer, 2020. [6](#)
- [58] Zhong-Qiu Wang and Ivan Tashev. Learning utterance-level representations for speech emotion and age/gender recognition using deep neural networks. In *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 5150–5154. IEEE, 2017. [2](#)
- [59] Yandong Wen, Bhiksha Raj, and Rita Singh. Face reconstruction from voice using generative adversarial networks. *Advances in neural information processing systems*, 32, 2019. [2](#), [3](#)
- [60] Cho-Ying Wu, Chin-Cheng Hsu, and Ulrich Neumann. Cross-modal perceptionist: Can face geometry be gleaned from voices? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10452–10461, 2022. [3](#), [6](#), [7](#)
- [61] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7623–7633, 2023. [2](#), [3](#), [4](#)
- [62] Zhen Xing, Qi Dai, Han Hu, Zuxuan Wu, and Yu-Gang Jiang. Simda: Simple diffusion adapter for efficient video generation. *arXiv preprint arXiv:2308.09710*, 2023. [3](#)
- [63] Chao Xu, Junwei Zhu, Jiangning Zhang, Yue Han, Wenqing Chu, Ying Tai, Chengjie Wang, Zhifeng Xie, and Yong Liu. High-fidelity generalized emotional talking face generation with multi-modal emotion space learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6609–6619, 2023. [1](#), [2](#)
- [64] Chao Xu, Shaoting Zhu, Junwei Zhu, Tianxin Huang, Jiangning Zhang, Ying Tai, and Yong Liu. Multimodal-driven talking face generation, face swapping, diffusion model. *arXiv preprint arXiv:2305.02594*, 2023. [2](#)
- [65] Fei Yin, Yong Zhang, Xiaodong Cun, Mingdeng Cao, Yanbo Fan, Xuan Wang, Qingyan Bai, Baoyuan Wu, Jue Wang, and Yujia Yang. Styleheat: One-shot high-resolution editable talking face generation via pretrained stylegan. *arXiv preprint arXiv:2203.04036*, 2022. [2](#)
- [66] Tao Yu, Runseng Feng, Ruoyu Feng, Jinming Liu, Xin Jin, Wenjun Zeng, and Zhibo Chen. Inpaint anything: Segment anything meets image inpainting. *arXiv preprint arXiv:2304.06790*, 2023. [5](#)
- [67] Wenxuan Zhang, Xiaodong Cun, Xuan Wang, Yong Zhang, Xi Shen, Yu Guo, Ying Shan, and Fei Wang. Sadtalker: Learning realistic 3d motion coefficients for stylized audio-driven single image talking face animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8652–8661, 2023. [1](#), [2](#), [5](#), [7](#)
- [68] Yabo Zhang, Yuxiang Wei, Dongsheng Jiang, Xiaopeng Zhang, Wangmeng Zuo, and Qi Tian. Controlvideo: Training-free controllable text-to-video generation. *arXiv preprint arXiv:2305.13077*, 2023. [3](#)
- [69] Zixing Zhang, Bingwen Wu, and Björn Schuller. Attention-augmented end-to-end multi-task learning for emotion prediction from speech. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6705–6709. IEEE, 2019. [2](#)
- [70] Zhimeng Zhang, Lincheng Li, Yu Ding, and Changjie Fan. Flow-guided one-shot talking face generation with a high-resolution audio-visual dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3661–3670, 2021. [6](#)
- [71] Weizhi Zhong, Chaowei Fang, Yinqi Cai, Pengxu Wei, Gangming Zhao, Liang Lin, and Guanbin Li. Identity-preserving talking face generation with landmark and appearance priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2023. [2](#)
- [72] Hang Zhou, Yu Liu, Ziwei Liu, Ping Luo, and Xiaogang Wang. Talking face generation by adversarially disentangled audio-visual representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9299–9306, 2019. [2](#)
- [73] Hang Zhou, Yasheng Sun, Wayne Wu, Chen Change Loy, Xiaogang Wang, and Ziwei Liu. Pose-controllable talking face generation by implicitly modularized audio-visual representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4176–4186, 2021. [2](#), [5](#), [7](#)
- [74] Hao Zhu, Wayne Wu, Wentao Zhu, Liming Jiang, Siwei Tang, Li Zhang, Ziwei Liu, and Chen Change Loy. Celebv-hq: A large-scale video facial attributes dataset. In *European conference on computer vision*, pages 650–667. Springer, 2022. [7](#)