

Efficient Emotional Adaptation for Audio-Driven Talking-Head Generation

Yuan Gan^{1,2}, Zongxin Yang^{1,2}, Xihang Yue^{1,2}, Lingyun Sun², Yi Yang^{1,2*}
¹ReLER, CCAI, Zhejiang University, China
²College of Computer Science and Technology, Zhejiang University, China

Abstract

Audio-driven talking-head synthesis is a popular research topic for virtual human-related applications. However, the inflexibility and inefficiency of existing methods, which necessitate expensive end-to-end training to transfer emotions from guidance videos to talking-head predictions, are significant limitations. In this work, we propose the Emotional Adaptation for Audio-driven Talking-head (EAT) method, which transforms emotion-agnostic talking-head models into emotion-controllable ones in a cost-effective and efficient manner through parameter-efficient adaptations. Our approach utilizes a pretrained emotion-agnostic talking-head transformer and introduces three lightweight adaptations (the Deep Emotional Prompts, Emotional Deformation Network, and Emotional Adaptation Module) from different perspectives to enable precise and realistic emotion controls. Our experiments demonstrate that our approach achieves state-of-the-art performance on widely-used benchmarks, including LRW and MEAD. Additionally, our parameter-efficient adaptations exhibit remarkable generalization ability, even in scenarios where emotional training videos are scarce or nonexistent. Project website: <https://yuangan.github.io/eat/>

1. Introduction

Recently, there has been increasing attention on synthesizing realistic talking heads due to their wide-ranging applications in industry, such as digital human animation [20, 22, 55], visual dubbing [43], and video content creation [50]. Audio-driven talking-head generation aims to produce realistic talking-head videos synchronized with speech. However, unlike speech, humans convey intentions through emotional expressions. Therefore, generating emotional talking heads is important to improve the fidelity of talking heads for real-world applications. To address this open problem, various forms of knowledge (such as human head models, emotions, audio, and vision) must be consid-

*Corresponding author.

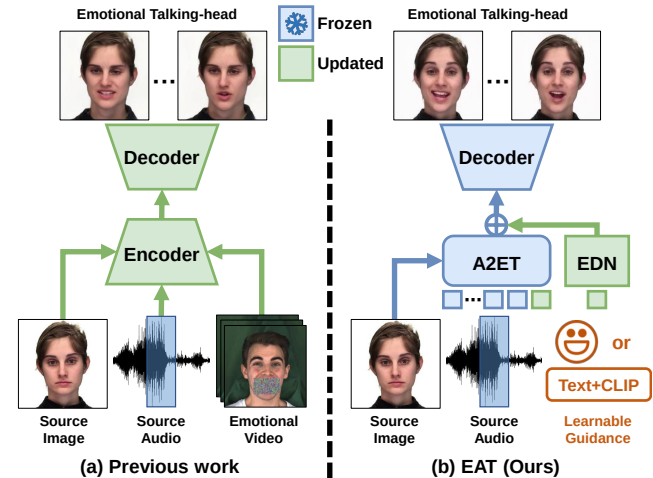


Figure 1. **Efficient emotional talking-head generation.** (a) Previous work trains or finetunes the whole network with augmented emotional driving videos. (b) Our EAT transforms emotion-agnostic talking-head models into emotion-controllable ones through flexible guidance, including emotional prompts or text-guided CLIP [45] supervision, by lightweight adaptations.

ered in constructing multi-knowledge representations [57].

Previous one-shot talking-head generation methods [67, 43, 55] focus on achieving audio-visual synchronization for emotion-agnostic talking-heads, which is a special case of realistic talking-heads. More recent works [23, 22, 34] pay attention to generating emotion-aware talking-heads. GC-AVT [34] and EAMM [22] are two methods that generate emotional videos using driven emotional videos and pose-guiding videos. GC-AVT [34] achieves explicit control over the expression, speech content, and pose of talking heads through granular pre-processing design. EAMM [22] synthesizes one-shot emotional talking-heads by adding augmented emotional source videos. Since the driven videos can introduce semantic ambiguity to mouth shapes, GC-AVT replaces the mouth part with the neighbor frame, while EAMM ignores the mouth part of driven emotional videos by using data augmentation. Additionally, these methods require training or finetuning the entire network at high costs for emotional talking-heads generation.

Although emotion-aware methods have made progress in the one-shot talking-head generation, they lack in-depth thinking in two key aspects. **(1) Architecture efficiency.** As a sub-task of talking-head generation, it is parameter-inefficient to train or finetune the entire emotional talking-head generation network. Furthermore, since large-scale emotion-agnostic talking-head data is more readily available than emotional data, it is worthwhile to consider how to efficiently reuse the knowledge learned from emotion-agnostic data. **(2) Guidance flexibility.** Previous methods prefer to transfer the driving videos to the target talking heads rather than directly learning the emotional representation. In practice, finding appropriate driven emotional videos requires taking into account factors such as resolution, occlusion, and even the length of the driven emotional videos and audio. Furthermore, prior research has neglected to consider lip shapes, which can result in unrealistic emotional expressions. For instance, according to FACS [18, 27], depressed lip corners are one of the key components of the sad expression.

To address the above limitations, a desirable approach should enable an efficient and flexible transfer of pretrained talking-head models to emotional talking-head generation tasks with lightweight emotional guidance, as illustrated in Fig. 1. There are two key advantages. Firstly, with the reused knowledge, we can readily and effortlessly apply the talking-head model to emotional talking-head generation tasks. Secondly, obtaining lightweight guidance is simpler and more adaptable in practical scenarios, such as text-guided zero-shot expression editing.

To realize the aforementioned paradigm, we propose an efficient Emotional Adaptation framework for audio-driven Talking-head (EAT) generation, which involves two stages. In the first stage, we enhance the unsupervised 3D latent keypoints representation [56] to capture emotional expressions. Then, we introduce the Audio-to-Expression Transformer (A2ET), which learns to map audio to enhanced 3D latent keypoints using large-scale talking-head datasets. In the second stage, we propose learnable guidance and adaptation modules for steering emotional expression generation. These include Deep Emotional Prompts for parameter-efficient emotional adaptation, a lightweight Emotional Deformation Network (EDN) for learning the emotional deformation of facial latent representation, and a plug-and-play Emotional Adaptation Module (EAM) for enhancing visual quality. Our approach enables rapid transfer of traditional talking-head models to emotional generation tasks with high-quality results and supports zero-shot expression editing with image-text models [45].

We conduct extensive experiments to assess the effectiveness of EAT on emotional talking-head generation. Compared to baseline competitors, EAT achieves superior performance without guiding emotional videos. Moreover,

based on the pretrained talking-head model, we can attain state-of-the-art (SOTA) performance in 2 hours with only 25% training data. The results indicate that our method is capable of generating more realistic talking-head videos. And with only text descriptions of emotions, we can achieve zero-shot talking-head editing.

In summary, the main contributions of our work are listed below:

- Our study introduces a new two-stage paradigm, called EAT, for addressing emotional talking-head tasks. Our experiments demonstrate that this paradigm outperforms previous methods with respect to both emotion manipulation and video quality in one-shot talking-head generation tasks.
- Our proposed architecture includes deep emotional prompts, an emotional deformation network, and an emotional adaptation module. This design enables the efficient transfer from generating talking heads without emotional expression to generating talking heads with emotional expression.
- To the best of our knowledge, our study is the first to introduce flexible guidance for talking-head adaptation. By utilizing image-text models, we can achieve zero-shot expression editing of talking-head videos, surpassing the capabilities of previous methods.

2. Related Work

Previous works have achieved great performance in audio-visual synchronization. However, there remain some challenges such as efficient knowledge transfer based on a large-scale pretrained talking-head model.

Audio-driven Talking Head Generation. Audio-driven talking-head generation [10, 51, 63, 7, 67, 20, 55, 64] based on deep learning has attracted lots of attention in recent years. Chen *et al.* [7] design a two-stage structure that leverages facial landmarks as an intermediate representation. MakeItTalk [67] generates one-shot talking heads based on disentangled speech and speaker. PC-AVS [64] generates arbitrary talking heads with pose control.

Following the groundbreaking work of transformers [52] in NLP [15, 46], recent works have achieved remarkable progress in video-related tasks, including action recognition [2, 19, 36], video segmentation [58, 59, 68, 41], cross-model understanding [32, 8], etc. In the talking-head field, AVCT [55] designs an audio-visual correlation transformer for generating talking-head videos. In this work, we enhanced 3D latent keypoints and applied transformers to generate more realistic talking heads.

Emotion-aware Talking Head Generation. Emotional talking-head generation has been studied recently for realistic talking-head generation. Pumarola *et al.* [44] introduce

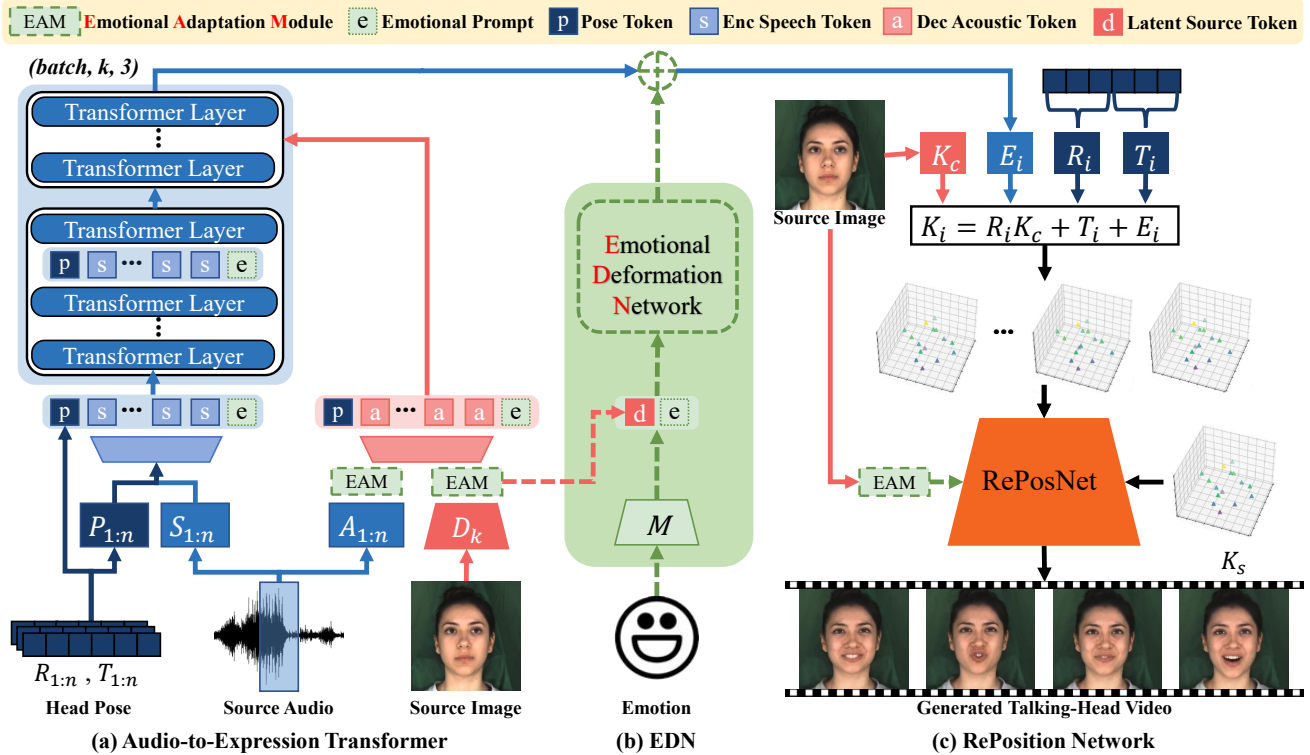


Figure 2. **Overview of EAT model.** (a) In the first stage, the Audio-to-Expression Transformer (A2ET) transfers latent source image representation, source audio and head pose sequences to 3D expression deformation. (b) In the second stage, the emotional guidance is injected into A2ET, Emotional Deformation Network (EDN) and Emotional Adaptation Module (EAM) for emotional talking-head generation, presented in dashed lines. (c) The RePos-Net takes the 3D source keypoints K_s and driven keypoints K_i to generate frames.

an unsupervised framework to generate facial videos with a specific expression. EVP [23] proposes emotional video portraits to produce more vivid results based on source video. However, the one-shot emotional talking-head generation has emerged recently. Sinha [49] generates pose-fixed emotional talking-heads with graph convolution. GC-AVT [34] train an emotion and pose controllable model with a granular pre-processing design. EAMM [22] synthesizes one-shot emotional talking heads with augmented emotional source videos. In our work, we achieve emotion control with efficient adaptation based on a pretrained talking-head model.

Efficient Finetuning. Efficient finetuning has been studied for knowledge transfer and many techniques have been proposed including residual adapter [47], bias tuning [5] and side tuning [61]. Recently, prompt [35] has attracted more attention in vision task. Inspired by prompt tuning in the language model, prompt tuning recently has been proposed in various visual tasks [45, 24, 66, 25, 65, 17, 33] for effectiveness and efficiency. CoOp [66] and VPT [25] utilize learnable prompt vectors and achieve better performance. Moreover, to increase generalizability, Co-CoOp [65] designed a lightweight network to learn prompt vectors for each image. In our work, we introduce deep

emotional prompts, an emotional deformation network and emotional adaptation modules to achieve efficient and effective emotion-related knowledge transfer.

3. Method

To avoid the expensive end-to-end training and finetuning of previous methods [22, 34], we propose a two-stage paradigm, the efficient Emotional Adaptation for audio-driven Talking-head (EAT) generation method. Firstly, we introduce the enhanced 3D latent representations and the emotion-agnostic talking-head pretraining using the Audio-to-Expression Transformer (A2ET). (Sec. 3.1) Secondly, we present a parameter-efficient emotional adaptation approach to quickly adapt the pretrained talking-head model for emotional talking-head tasks. This approach includes deep emotional prompts, an Emotional Deformation Network (EDN), and an Emotional Adaptation Module (EAM). (Sec. 3.2) At last, we introduce our training objectives in detail. (Sec. 3.3)

3.1. Emotion-agnostic Pretraining

Compared to inefficient end-to-end training for emotional talking-head [34, 22], we explore leveraging the knowledge gained from an adaptable model that has been

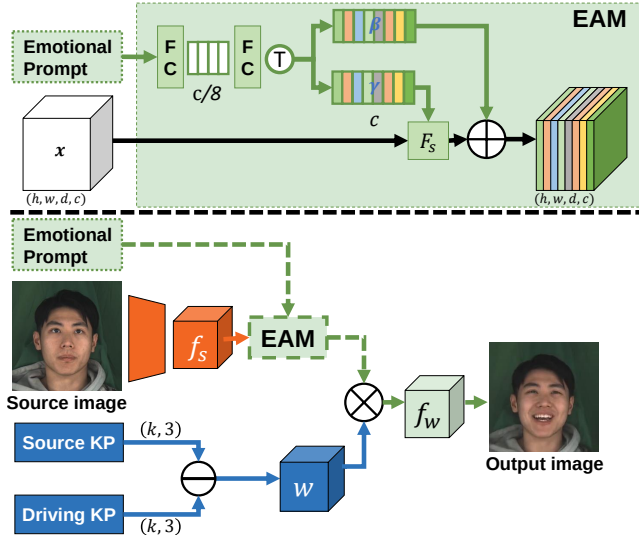


Figure 3. **The architecture of RePos-Net and EAM.** RePos-Net extracts 3D appearance features f_s from the source image. Given the 3D source keypoints and driving keypoints, RePos-Net predicts the 3D flow warp matrix w to transform the 3D feature f_s and generates the output frame. EAM transfers emotional guidance to emotion-conditioned features with the learned γ and β .

pre-trained on emotion-agnostic datasets. The model allows for rapid adaptations to downstream tasks, including emotional talking-head generation. To achieve such an emotion-agnostic pertaining model, we firstly enhance the 3D latent representation [56] to capture the subtle expression better. Then an Audio-to-Expression Transformer (A2ET) is introduced to capture the temporal context of audio clips, generate audio-visual synchronized 3D latent expression sequences for talking-head generation, and support efficient emotional adaptations (Sec. 3.3).

3.1.1 Enhanced Latent Representation.

Given a talking-head frame i , the 3D latent keypoints K_i , which are learned unsupervisedly, are composed of four components: identity-specific canonical keypoints K_c , rotation matrix R_i , translation T_i , and expression deformation E_i . These components are then combined with the following equation:

$$K_i = R_i K_c + T_i + E_i. \quad (1)$$

Based on 3D latent keypoints, *RePosition Network* (RePos-Net) [56] can transfer facial expressions from one person to another, as illustrated in Fig. 3. However, we observe that this kind of transfer fails to account for other facial expression elements, such as eyebrows, lip corners, etc. Therefore, we enhance the latent representation for more realistic talking-head generation with these modifications:

- We remove the deformation prior loss in OSFV [56]

which penalizes the magnitude of the keypoints deformation. This allows our latent keypoints to capture more subtle changes in facial expressions.

- We use the MEAD dataset [54] to acquire labeled and paired face data from neutral and emotional videos of the same identity. This helps the network learn more expressive faces from expression changes.
- To avoid the influence of expression-irrelevant background, we only compute losses on the facial part. And we augment our paired data with the affectnet [40] dataset to improve generalizability.

These modifications enhance the representation ability of the learned 3D latent keypoints, which are the objective of our A2ET model.

3.1.2 Audio-to-Expression Transformer.

Since the 3D latent keypoints are specific to the source identity and more complex than 2D latent keypoints [48, 55, 22], predicting the 3D keypoints sequence directly is a challenging task. We notice that the facial expressions are mainly represented by the expression deformation E_i in 3D latent keypoints. Thus, the purpose of A2ET is to learn the audio-visual synchronized expression deformation, which is composed of audio-visual feature extraction and expression deformation prediction.

Audio-visual Feature Extraction. Previous work [55] generates emotion-agnostic talking heads with transformer and phonemes. However, training transformer requires a large dataset, and phoneme extraction is challenging in noisy or accented speech. To address these limitations, we train our A2ET model on a large dataset Voxceleb2 [11] and extract speech features $S_{1:n}$ and acoustic features $A_{1:n}$ as inputs. The audio semantic features $S_{1:n}$ are extracted via a speech recognition model [1] from the MFCC features. To derive acoustic features $A_{1:n}$, we design an audio encoder to encode the mel spectrogram extracted with 80 mel bins and 1025 frequency bins.

As shown in Fig. 2(a), given frame i , we extract its semantic context features from $2w + 1$ audio frames. Initially, speech features $S_{i-w:i+w}$ and head pose features $P_{i-w:i+w}$ are converted into speech tokens. The 6DoF of frame i is encoded into a pose token p . The A2ET encoder takes these tokens as input. Subsequently, to capture subtle mouth movements, we encode acoustic features $A_{i-w:i+w}$ and the latent source image representation with an audio encoder [31] and keypoint detector D_k . These representations are fused to obtain acoustic tokens that are used by the A2ET decoder to output the feature of $2w + 1$ tokens.

Expression Deformation Prediction. The expression deformation E_i , composed of k 3D offsets, can be predicted from the feature of central frame i , but optimizing the 3D




Figure 4. **The qualitative results of the one-shot emotional talking-head generation cases.** Results of *happy* and *surprised* are on the left and right. The top row shows the identity and driving emotion. The second row shows the content ground-truth frames. The left source face from LRW [12]. The right face from CREMA-D [6].

expression motions leads to convergence issues. We observed that the 3D keypoints learned by self-supervision exhibit inherent inter-dependencies, with only a few keypoints influencing facial expressions. To address this, the principal component analysis (PCA) of E_i is adopted to reduce the dimensionality and eliminate nonessential information. Then we can predict the 3D expression deformation from the audio features.

3.2. Efficient Emotional Adaptation

Traditional audio-driven talking-head methods [7, 67, 20, 55, 64] have made impressive advancements in emotion-agnostic talking-head generation tasks. However, to achieve realistic talking heads, emotional expression is crucial. Hence, we present a novel approach that includes three parameter-efficient modules for swift emotional adaptation from emotion-agnostic models. These modules comprise Deep Emotional Prompts, the Emotional Deformation Network (EDN), and the Emotional Adaptation Module (EAM), specifically designed to enable efficient emotional adaptation of pre-trained A2ET. Our approach allows for lightweight adaptations, which offer flexibility in guiding downstream tasks, such as zero-shot expression editing. **Emotional Guidance.** One straightforward idea is to generate emotional talking heads using learnable guidance conditioned on emotions. We posit that each emotional type belongs to a distinct sub-domain in latent space. As shown in Fig. 2 (b), a mapping network M is adopted to extract emotion-conditional guidance with a latent code $z \in \mathcal{U}^{16}$. the latent code is sampled from a Gaussian distribution, which is commonly used in generative models [29, 9]. This emotional guidance is applied to steer the generation of emotional expressions.

Deep Emotional Prompt. To achieve parameter-efficient

emotional adaptation, we include the guidance as an additional input token of the A2ET transformer layer, as shown in Fig. 2(a). We separately introduce shallow and deep emotional prompts into the A2ET transformer architecture, with the shallow prompt added to the first layer and the deep prompt added to every layer thereafter. Our results in Table 4 demonstrate that the deep prompt leads to better emotional expression transfer compared to the shallow prompt. However, we also observe that incorporating emotional prompts can have a detrimental effect on audio-visual synchronization. Generating emotional expressions using fixed transformer weights while ensuring audio-visual synchronization may pose a challenge for the prompts.

Emotional Deformation Network. We observe that the decoupled 3D implicit representations in Eq. 1 exhibit linear additivity. Furthermore, emotional talking heads show emotional deformation that is not present in traditional talking-heads. To complement E_i , one intuitive approach is to include an emotional expression deformation term:

$$E'_i = E_i + \Delta E_i, \quad (2)$$

where E'_i represents emotional expression deformation, E_i represents speech-related expression deformation predicted by A2ET, and ΔE_i represents emotion-related expression deformation. To predict ΔE_i , we design a sub-network called the Emotional Deformation Network (EDN), which is depicted in Fig. 2(b). EDN utilizes the A2ET encoder architecture to predict ΔE_i with emotional guidance and the source latent representation token. To accelerate adaptation, we initialized EDN with the pretrained A2ET encoder. To update E_i with E'_i , we can get emotional 3D latent keypoints using Eq. 1.

Emotional Adaptation Module. To enhance the visual quality, we have designed a lightweight, plug-and-play

	LRW [12]				MEAD [54]				
	PSNR/SSIM↑	FID↓	SyncNet↑	M/F-LMD↓	PSNR/SSIM↑	FID↓	SyncNet↑	M/F-LMD↓	Acc _{emo} ↑
ATVG [7]	18.40/0.64	51.56	2.73	2.69/3.31	17.64/0.56	99.42	1.80	2.77/3.74	17.36
Wav2Lip [43]	22.80/0.73	7.44	7.59	1.58 /2.47	19.12/0.57	67.49	8.97	3.11/3.71	17.87
MakeItTalk [67]	21.67/0.69	3.37	3.28	2.16/2.99	18.79/0.55	51.88	5.28	3.61/4.00	15.23
AVCT [55]	21.72/0.68	2.01	4.63	2.55/3.23	18.43/0.54	39.18	6.02	3.82/4.33	15.64
PC-AVS [64]	23.32/0.72	4.64	7.36	1.54 /2.11	20.60 /0.61	53.04	8.60	2.66/2.70	11.88
EAMM [22]	22.34/0.71	6.44	4.67	1.81/2.37	20.55/0.66	22.38	6.62	2.19 /2.55	49.85
Pretrain (Ours)	23.97 /0.76	1.89	6.30	1.95/2.12	20.32/0.61	26.71	8.09	2.83/2.99	25.18
EAT (Ours)	24.11 /0.77	3.52	6.22	1.79/2.08	21.75 /0.68	19.69	8.28	2.25 /2.47	75.43
Ground Truth	∞ /1.00	0	7.06	0.00/0.00	∞ /1.00	0	7.76	0.00/0.00	84.37

Table 1. **Quantitative comparisons with state-of-the-art methods on LRW [12] and MEAD [54].** We present the results of pretrained A2ET and full EAT model on both LRW and MEAD. M/F-LMD denotes the landmark distance of mouth and face. "↑": higher is better. "↓": lower is better. **Red**: the 1st score. **Blue**: the 2nd score.

Method	Wav2Lip	PC-AVS	EAMM	EAT	GT
Lip-sync	3.86	3.90	3.64	3.99	4.59
Quality	2.69	3.19	2.89	3.35	4.59
Acc _{emo}	13%	20%	35%	50%	66%

Table 2. **User Study on CREMA-D and LRW.** Lip-sync and Quality represent the audio-visual synchronization and visual quality. Emotion classification accuracy (Acc_{emo}) evaluates the effectiveness of methods in producing emotional expressions.

adaptation module called the Emotional Adaptation Module (EAM) that can generate emotion-conditioned features. As shown in Fig. 3, the module takes in the guidance embedding e and processes it through two fully connected (FC) layers to obtain a set of channel weights γ and bias β . And we use the tanh activation function to confine the γ and β values to the range $[-1, 1]$:

$$\gamma, \beta = \tanh(\text{FC}(\text{ReLU}(\text{FC}(e))))). \quad (3)$$

Once we have obtained γ and β , we can input the feature x to obtain the emotional feature, which is calculated using the following equation:

$$EAM(x) = F_s(1 + \gamma, x) + \beta, \quad (4)$$

where F_s denotes the channel-wise multiplication. As shown in Figure 2, the EAM can be inserted into RePosNet, as well as the audio and image feature extractors.

Zero-shot Expression Editing. Owing to the quick adaptation capabilities of our EAT, we can achieve zero-shot text-guided expression editing of talking heads by distilling knowledge from CLIP [45], a large-scale vision-language pertaining model. This unique ability sets our work apart from the latest research [22], as it eliminates the need for emotional training data and enables generalization to applications that require rare expressions.

Specifically, our goal is to employ CLIP loss to learn emotional guidance correlated with text-described expres-

sions. To achieve this, we extract the head poses, source audio and the first frame from the target video as the input. Besides, a target expression description is taken for finetuning. Utilizing the refined EAT model and our training loss, we add an *additional* CLIP loss [42] to finetune the mapping network and EAM module only. In detail, we extract image embeddings from the predicted talking face using the image encoder of CLIP, and text embeddings from the description using its text encoder. We then iteratively optimize the distance between the image and text embeddings to align the generated talking face with the input text.

3.3. Training Objectives

For supervised learning, the loss is calculated as follows:

$$\mathcal{L} = \lambda_{lat}\mathcal{L}_{lat} + \lambda_{sync}\mathcal{L}_{sync} + \lambda_{rec}\mathcal{L}_{rec}, \quad (5)$$

where λ_{lat} , λ_{sync} and λ_{rec} are hyper-parameters that re-weight the corresponding term. As for zero-shot editing, we replace the $\lambda_{rec}\mathcal{L}_{rec}$ with the CLIP loss since there is no ground-truth video. In the following, we will discuss each training loss in detail.

Latent Loss. Latent loss is applied to optimize the predicted latent keypoints:

$$\mathcal{L}_{lat} = \frac{1}{N} \sum_{i=1}^N (\|PE_i - \hat{PE}_i\|_2^2 + \|K_i - \hat{K}_i\|_2^2), \quad (6)$$

where N denotes the frame length of a sampled audio clip in each batch. PE_i represents the predicted PCA of expression deformation in i frame. K_i is the transferred 3D latent keypoints according to Eq. 1. \hat{PE}_i and \hat{K}_i are the corresponding ground-truth of the frame i . As the emotional expression deformation ΔE_i is not included in P_i , we only use the loss of 3D keypoints in Eq. 6 while training EDN.

Sync Loss. The synchronization loss is introduced in Wav2Lip [43]. Based on the structure of SyncNet [13], we

train an expert to discriminate the audio-visual synchronization in neutral and emotional datasets. For a sampled audio clip in each batch, we computed the synchronization loss of the generated video using the following equation:

$$\mathcal{L}_{\text{sync}} = -\log\left(\frac{v \cdot s}{\max(\|v\|_2 \cdot \|s\|_2, \epsilon)}\right). \quad (7)$$

The input speech embedding s and the generated video embedding v are extracted by the speech encoder and image encoder in SyncNet, respectively.

Reconstruction Loss. To improve the expression generation, we employ \mathcal{L}_1 reconstruction loss only in the facial region. Additionally, to generate sharper frames, we apply perceptual loss [26] to the entire frames using a pretrained VGG19 model.

CLIP Loss. The CLIP loss calculates the similarity between the embeddings of the generated face and the text description with a pretrained CLIP model. Specifically, the CLIP loss is calculated by taking the cosine similarity between the normalized embedding of the image and text.

4. Experiments

4.1. Experimental Setup

Implementation details. The videos are sampled to 25 FPS and the audio sample rate is 16KHz. The videos are cropped and resized to 256×256 . To synchronize the audio features and the video, We extract the mel-spectrogram [37] by configuring the window length and hop length to 640. The number of keypoints k used in EAT is 15. The mapping network of emotional prompts consists of shared four MLP layers and unshared four MLP layers for every kind of emotion. We enhance the 3D latent keypoints for 48 hours and pretrain the A2ET with enhanced latent keypoints for 48 hours. Then we finetune the EAT architecture for 6 hours. Our work is based on 4 NVIDIA 3090 GPUs.

Datasets. The training dataset consists of videos from VoxCeleb2 [11] and MEAD [54]. MEAD is a high-quality emotional talking-head video set with 8 kinds of emotions. To ensure fair comparisons, we split the MEAD dataset into training and testing sets based on identity, using the same test identities as EAMM [22]. To learn large head pose changes, we select about 8,000 emotional videos from the VoxCeleb2 [11] dataset with Emotion-FAN [39] for finetuning. To obtain the PCA of the enhanced 3D keypoints, we extract the largest 32 eigenvalue matrix and mean of 2,500 videos from the training set.

Baselines. We compare with SOTA one-shot talking-head generation methods on LRW [12] and MEAD [54] test set. They are ATVG [7], Wav2Lip [43], MakeItTalk [67], AVCT [55], PC-AVS [64] and EAMM [22].

Metric. We assess the quality of synthesized emotional videos with the following metrics:

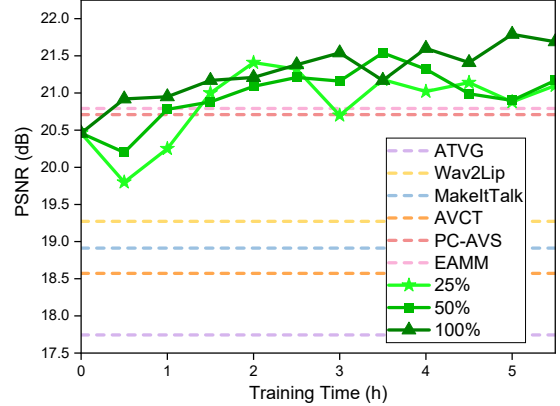


Figure 5. **Tuning Efficiency.** Our EAT exhibits exceptional tuning efficiency, achieving state-of-the-art performance with just 50% MEAD training data in a one-hour fine-tuning session, or even with only 25% data in a two-hour session.




Figure 6. The percentage of parameters in EAT.

Image quality. We utilize PSNR, SSIM and Fréchet Inception Distance score (FID) [21] to measure the image quality of synthesized videos.

Audio-visual synchronization. We evaluate the audio-visual synchronization of the synthesized videos using the confidence score of SyncNet [13]. In addition, the distance between the landmarks of the mouth (M-LMD) [7] is used to indicate speech content consistency, while the distance between the landmarks of the whole face (F-LMD) represents the accuracy of the pose and expressions.

Emotional accuracy. To assess the emotional accuracy (Acc_{emo}) of the generated emotions, we fine-tune the Emotion-Fan [39] using the training set of MEAD.

4.2. Talking-head Generation

To verify the effectiveness of EAT, we do experiments on emotion-agnostic and emotional talking-head generation.

Emotion-agnostic Talking-head Generation. For one-shot emotion-agnostic talking-head generation, we test on the LRW test set, which comprises 25k neutral videos. We take the first frame as the source image for every test video. As presented in Table 1, our method outperforms other methods in terms of visual quality for emotion-agnostic talking-head generation. Furthermore, our EAT can im-



Figure 7. **CLIP-based zero-shot editing.** The expressions are provided by “*text description*”. The neutral videos and source images are in-the-wild from LRW [12].



Figure 8. Additional zero-shot results of our EAT. Emotional expressions are given by “*text description*”. The first row of each text shows the editing results and the second row shows the generated talking head of a different identity with the learned guidance. Please refer to our video for more details. Source images are from MEAD[54] and LRW[12].

prove the performance of the pretrained talking-head model. Wav2Lip and PC-AVS risk overfitting the pretrained lip synchronization scoring model since their sync scores surpass the ground truth. Additionally, Wav2Lip only generates mouth parts without facial expression and head pose.

Emotional Talking-head Generation. We follow the setting of EAMM to compare the emotional talking-head generation on the public-available MEAD test set. For all methods, the neutral source frames are from EAMM.

Table 1 shows our EAT can achieve the best performance

Method	PSNR \uparrow	M/F-LMD \downarrow	Sync \uparrow	Acc $_{emo}$ \uparrow
OSFV [56]	22.39	1.60/2.12	6.68	30
Enhanced	24.49	1.09/1.49	7.49	86

Table 3. **Ablation study of enhanced latent representation.** To validate the effectiveness of our enhanced latent representation, we generate emotional videos driven by source images and emotional videos in the MEAD test set.

Prompt	PSNR \uparrow	M/F-LMD \downarrow	Sync \uparrow	Acc $_{emo}$ \uparrow
w/o	20.46	2.85/2.99	8.12	25
Shallow	21.19	2.50/2.63	7.63	57
Deep	21.23	2.36/2.48	7.83	84

Table 4. **Ablation study of prompt.** To verify the effectiveness of shallow and deep emotional prompts, we produce videos using source images and emotional videos from the MEAD test set.

in most metrics. Specifically, EAT achieves better video quality and higher emotion accuracy than other approaches. These findings support the superiority of the emotional representation learned by our proposed method. Fig. 4 visually demonstrates our capability to produce authentic and comprehensive emotional expressions in the one-shot setting. Note that AVCT [55] is not capable of controlling pose explicitly even with ground-truth 6DoF. For more results, please refer to our supplementary.

4.3. User Study

To evaluate the generated emotional talking heads with in-the-wild images, we conducted a user study with 14 participants to assess lip-sync, video quality, and emotion classification. The maximum score value is 5. To ensure a diverse set of images, we randomly sampled 16 images from both the CREMA-D and LRW datasets, which are not included in the training data. Additionally, we used audio from MEAD test set to generate a total of 32 videos (4×8 emotions) for each method. As shown in Table 2, our method achieved the best scores for lip-sync, video quality, and emotion accuracy, validating the effectiveness of our EAT architecture in generating emotional talking heads. Besides, the results show that the sync value in Table 1 is inaccurate for emotional talking-head due to the SyncNet model being trained only with neutral talking-head videos. More in-the-wild results can be found in the supplementary.

4.4. Tuning Efficiency

During the finetuning of the second stage, we conducted periodic tests on a subset of the MEAD test set every half hour. This enabled us to demonstrate the time and data efficiency of our EAT. EAT can efficiently adapt the pretrained A2ET model to sub-tasks, even with limited data. As shown in Fig. 5, EAT can surpass the SOTA results within one hour

Ablation	(A)	(B)	(C)	EAT
Prompt		✓	✓	✓
EDN			✓	✓
EAM				✓
PSNR↑	20.46	21.23	21.40	21.79
M/F-LMD↓	2.85/2.99	2.36/2.48	2.28/ 2.41	2.22/2.43
Sync↑	8.12	7.83	7.83	8.22
Acc _{emo} ↑	25	84	81	67

Table 5. **Ablation study of each component.** Each component contributes to the improvement of the video quality, thus verifying its effectiveness.

	✓	✓	✓	✓
\mathcal{L}_{per}	✓	✓	✓	✓
\mathcal{L}_{lat}		✓	✓	✓
\mathcal{L}_{sync}			✓	✓
\mathcal{L}_1				✓
PSNR↑	21.52	21.61	21.31	21.79
Sync↑	5.50	5.66	8.13	8.22

Table 6. **Ablation study of each loss.** Each loss contributes to the improvement of the video quality or synchronization value.

with full or half data. We can also achieve comparable performance with just a quarter of the data within two hours. Moreover, as shown in Fig 6 the emotional adaptation modules require only 6.85% additional parameters compared to the pretrained model. Deep emotional prompts account for 5.30%, EDN for 1.17%, and EAM for 0.38% of the additional parameters. These results demonstrate the effectiveness and efficiency of our EAT approach.

4.5. Zero-shot Expression Editing

We conduct zero-shot expression editing with the CLIP [45] model to generate a novel emotional talking head video, as shown in Fig. 7. Given a neutral video, we treat the first frame as the source image and edit the expression by text descriptions. Based on EAT architecture, we learn the emotional guidance and EAM with an additional CLIP loss [45]. We notice that text descriptions will determine the editing performance and need careful design. Besides, the learned latent code from one video and guidance texts can also be applied to another video. As shown in Fig. 8, We present different identity results manipulated by a learned latent code.

4.6. Ablation Study

To assess the effectiveness and importance of various aspects of EAT, We conduct several ablation studies on our proposed architecture and modules.

Enhanced Latent Representation. We compare the latent representation before and after our proposed enhancement. As shown in Table 3, the superior performance

in face reenactment shows that our pre-trained model can capture a wider range of emotional facial movements than OSFV [56], which has been discussed in Sec. 3.1.1.

Prompt. To study the effect of different types of prompts, we conduct experiments with shallow and deep emotional prompts. Table 4 shows that deep prompt learns emotional expression deformation better than shallow prompt, although it has a side effect on synchronization.

Each Component. To verify the effectiveness of our proposed modules, we conduct ablation experiments by removing one component every time. Table 5 shows that all three components can improve video quality. Deep emotional prompts can transfer the talking-head knowledge to emotional talking-head generation at the expense of synchronization. Although deep emotional prompts lead to intense emotional expressions, the outputs deviate from the ground truth. By incorporating EDM and EAM, the image quality and fidelity toward the ground truth are enhanced, as evidenced by the rise in PSNR/SSIM values. Yet, this comes at the expense of emotion intensity and accuracy. Please refer to our supplementary for visual analysis.

Each Loss. As shown in Table 6, we conduct an ablation study on perceptual loss, latent loss, sync loss, and \mathcal{L}_1 loss in the finetuning stage. It demonstrates that sync loss contributes to synchronization while others contribute to expression fidelity.

5. Conclusion

In this paper, we propose an efficient emotional adaptation paradigm for audio-driven talking-head generation, consisting of two stages. First, we enhance 3D latent representation and develop a transformer architecture A2ET to achieve emotion-agnostic talking-head generation. Second, we introduce learnable guidance for emotional expression control through our deep emotional prompts, EDN, and EAM. With these adaptation modules, EAT can quickly transfer the pretrained talking-head model to emotional talking-head generation. Experiments demonstrate that our EAT is the first parameter-efficient and effective paradigm for emotional talking-head generation.

Limitations and Broader Impact. 1) The drawbacks of emotional training data, such as the diversity of background and head poses, will affect the generalizability of our EAT. 2) Our approach paves the way for broader talking-head applications, including zero-shot or one-shot emotional talking-head generation.

Acknowledgements. This work is supported by the National Key R&D Program of China (2022YFB3303300) and the Fundamental Research Funds for the Central Universities (No. 226-2023-00048).

References

- [1] Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, et al. Deep speech 2: End-to-end speech recognition in english and mandarin. In *International conference on machine learning*, pages 173–182. PMLR, 2016. [4](#)
- [2] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *ICCV*, pages 6836–6846, 2021. [2](#)
- [3] Lisa Feldman Barrett, Ralph Adolphs, Stacy Marsella, Aleix M Martinez, and Seth D Pollak. Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psychological science in the public interest*, 20(1):1–68, 2019. [16](#)
- [4] Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette N Chang, Sungbok Lee, and Shrikanth S Narayanan. Iemocap: Interactive emotional dyadic motion capture database. *Language resources and evaluation*, 42:335–359, 2008. [16](#)
- [5] Han Cai, Chuang Gan, Ligeng Zhu, and Song Han. Tinytl: Reduce memory, not parameters for efficient on-device learning. *Advances in Neural Information Processing Systems*, 33:11285–11297, 2020. [3](#)
- [6] Houwei Cao, David G Cooper, Michael K Keutmann, Ruben C Gur, Ani Nenkova, and Ragini Verma. Crema-d: Crowd-sourced emotional multimodal actors dataset. *IEEE transactions on affective computing*, 5(4):377–390, 2014. [5](#), [13](#)
- [7] Lele Chen, Ross K Maddox, Zhiyao Duan, and Chenliang Xu. Hierarchical cross-modal talking face generation with dynamic pixel-wise loss. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7832–7841, 2019. [2](#), [5](#), [6](#), [7](#), [14](#)
- [8] Yangming Cheng, Liulei Li, Yuanyou Xu, Xiaodi Li, Zongxin Yang, Wenguan Wang, and Yi Yang. Segment and track anything. *arXiv preprint arXiv:2305.06558*, 2023. [2](#)
- [9] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8188–8197, 2020. [5](#)
- [10] Joon Son Chung, Amir Jamaludin, and Andrew Zisserman. You said that? *BMVC*, 2017. [2](#)
- [11] Joon Son Chung, Arsha Nagrani, and Andrew Zisserman. VoxCeleb2: Deep Speaker Recognition. In *Interspeech 2018*, pages 1086–1090. ISCA, Sept. 2018. [4](#), [7](#), [13](#)
- [12] Joon Son Chung and Andrew Zisserman. Lip reading in the wild. In *Asian conference on computer vision*, pages 87–103. Springer, 2016. [5](#), [6](#), [7](#), [8](#), [14](#), [16](#)
- [13] Joon Son Chung and Andrew Zisserman. Out of time: automated lip sync in the wild. In *Asian conference on computer vision*, pages 251–263. Springer, 2016. [6](#), [7](#)
- [14] Davide Cozzolino, Andreas Rossler, Justus Thies, Matthias Nießner, and Luisa Verdoliva. Id-reveal: Identity-aware deepfake video detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15108–15117, 2021. [16](#)
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, pages 4171–4186, 2019. [2](#)
- [16] Michail Christos Doukas, Evangelos Ververas, Viktoriia Sharmanska, and Stefanos Zafeiriou. Free-headgan: Neural talking head synthesis with explicit gaze control. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. [16](#)
- [17] Yu Du, Fangyun Wei, Zihe Zhang, Miaojing Shi, Yue Gao, and Guoqi Li. Learning to prompt for open-vocabulary object detection with vision-language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14084–14093, 2022. [3](#)
- [18] Paul Ekman and Wallace V Friesen. Facial action coding system. *Environmental Psychology & Nonverbal Behavior*, 1978. [2](#)
- [19] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *ICCV*, pages 6824–6835, 2021. [2](#)
- [20] Yudong Guo, Keyu Chen, Sen Liang, Yong-Jin Liu, Hujun Bao, and Juyong Zhang. Ad-nerf: Audio driven neural radiance fields for talking head synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5784–5794, 2021. [1](#), [2](#), [5](#)
- [21] 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. [7](#)
- [22] 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 *Special Interest Group on Computer Graphics and Interactive Techniques Conference Proceedings*, pages 1–10, Vancouver BC Canada, Aug. 2022. ACM. [1](#), [3](#), [4](#), [6](#), [7](#), [13](#), [14](#)
- [23] 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. [1](#), [3](#)
- [24] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021. [3](#)
- [25] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *European Conference on Computer Vision (ECCV)*, 2022. [3](#)
- [26] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. [7](#)

- [27] Dacher Keltner, Disa Sauter, Jessica Tracy, and Alan Cowen. Emotional expression: Advances in basic emotion theory. *Journal of nonverbal behavior*, 43(2):133–160, 2019. [2](#)
- [28] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015. [13](#)
- [29] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations*, 2014. [5](#)
- [30] Jean Kossaifi, Robert Walecki, Yannis Panagakis, Jie Shen, Maximilian Schmitt, Fabien Ringeval, Jing Han, Vedhas Pandit, Antoine Toisoul, Björn Schuller, et al. Sewa db: A rich database for audio-visual emotion and sentiment research in the wild. *IEEE transactions on pattern analysis and machine intelligence*, 43(3):1022–1040, 2019. [16](#)
- [31] Sangeun Kum and Juhan Nam. Joint Detection and Classification of Singing Voice Melody Using Convolutional Recurrent Neural Networks. *Applied Sciences*, 9(7):1324, Mar. 2019. [4](#)
- [32] Kexin Li, Zongxin Yang, Lei Chen, Yi Yang, and Jun Xiao. Catr: Combinatorial-dependence audio-queried transformer for audio-visual video segmentation. In *Proceedings of the 31th ACM International Conference on Multimedia*, 2023. [2](#)
- [33] Muheng Li, Lei Chen, Yueqi Duan, Zhilan Hu, Jianjiang Feng, Jie Zhou, and Jiwen Lu. Bridge-prompt: Towards ordinal action understanding in instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19880–19889, 2022. [3](#)
- [34] 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. [1, 3](#)
- [35] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing. *ACM Computing Surveys*, page 3560815, Sept. 2022. [3](#)
- [36] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *CVPR*, pages 3202–3211, 2022. [2](#)
- [37] Beth Logan. Mel frequency cepstral coefficients for music modeling. In *International Symposium on Music Information Retrieval*. Citeseer, 2000. [7](#)
- [38] Momina Masood, Marriam Nawaz, Khalid Mahmood Malik, Ali Javed, and Aun Irtaza. Deepfakes generation and detection: State-of-the-art, open challenges, countermeasures, and way forward. *arXiv preprint arXiv:2103.00484*, 2021. [16](#)
- [39] 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, Taipei, Taiwan, Sept. 2019. IEEE. [7](#)
- [40] A. Mollahosseini, B. Hasani, and M. H. Mahoor. Affectnet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Transactions on Affective Computing*, PP(99):1–1, 2017. [4, 13](#)
- [41] Xiao Pan, Peike Li, Zongxin Yang, Huiling Zhou, Chang Zhou, Hongxia Yang, Jingren Zhou, and Yi Yang. In-n-out generative learning for dense unsupervised video segmentation. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 1819–1827, 2022. [2](#)
- [42] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2085–2094, 2021. [6](#)
- [43] KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Namboodiri, and CV Jawahar. A lip sync expert is all you need for speech to lip generation in the wild. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 484–492, 2020. [1, 6, 7, 14, 15](#)
- [44] Albert Pumarola, Antonio Agudo, Aleix M Martinez, Alberto Sanfeliu, and Francesc Moreno-Noguer. Ganimation: Anatomically-aware facial animation from a single image. In *Proceedings of the European conference on computer vision (ECCV)*, pages 818–833, 2018. [2, 13](#)
- [45] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. [1, 2, 3, 6, 9](#)
- [46] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. [2](#)
- [47] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Learning multiple visual domains with residual adapters. *Advances in neural information processing systems*, 30, 2017. [3](#)
- [48] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. *Advances in Neural Information Processing Systems*, 32:7137–7147, 2019. [4](#)
- [49] Sanjana Sinha, Sandika Biswas, Ravindra Yadav, and Brojeshwar Bhowmick. Emotion-controllable generalized talking face generation. *International Joint Conferences on Artificial Intelligence Organization*, 2022. [3](#)
- [50] Linsen Song, Wayne Wu, Chaoyou Fu, Chen Qian, Chen Change Loy, and Ran He. Everything’s talkin’: Pareidolia face reenactment. *arXiv preprint arXiv:2104.03061*, 2021. [1](#)
- [51] Yang Song, Jingwen Zhu, Dawei Li, Xiaolong Wang, and Hairong Qi. Talking face generation by conditional recurrent adversarial network. *arXiv preprint arXiv:1804.04786*, 2018. [2](#)
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [2](#)
- [53] Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Realistic speech-driven facial animation with gans. *International Journal of Computer Vision*, 128:1398–1413, 2020. [16](#)

- [54] 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. [4](#), [6](#), [7](#), [8](#), [14](#), [15](#)
- [55] Suzhen Wang, Lincheng Li, Yu Ding, and Xin Yu. One-shot talking face generation from single-speaker audio-visual correlation learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2531–2539, 2022. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [8](#), [14](#)
- [56] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10039–10049, 2021. [2](#), [4](#), [8](#), [9](#), [13](#)
- [57] Yi Yang, Yueting Zhuang, and Yunhe Pan. Multiple knowledge representation for big data artificial intelligence: framework, applications, and case studies. *Frontiers of Information Technology & Electronic Engineering*, 22(12):1551–1558, 2021. [1](#)
- [58] Zongxin Yang, Yunchao Wei, and Yi Yang. Associating objects with transformers for video object segmentation. *Advances in Neural Information Processing Systems*, 34, 2021. [2](#)
- [59] Zongxin Yang and Yi Yang. Decoupling features in hierarchical propagation for video object segmentation. *Advances in Neural Information Processing Systems*, 2022. [2](#)
- [60] Stefanos Zafeiriou, Dimitrios Kollias, Mihalis A Nicolaou, Athanasios Papaioannou, Guoying Zhao, and Irene Kotzia. Aff-wild: Valence and arousal ‘in-the-wild’ challenge. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*, pages 1980–1987. IEEE, 2017. [16](#)
- [61] Jeffrey O Zhang, Alexander Sax, Amir Zamir, Leonidas Guibas, and Jitendra Malik. Side-tuning: a baseline for network adaptation via additive side networks. In *European Conference on Computer Vision*, pages 698–714. Springer, 2020. [3](#)
- [62] Yufeng Zheng, Victoria Fernández Abrevaya, Marcel C Bühler, Xu Chen, Michael J Black, and Otmar Hilliges. Im avatar: Implicit morphable head avatars from videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13545–13555, 2022. [16](#)
- [63] 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*, volume 33, pages 9299–9306, 2019. [2](#)
- [64] 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](#), [6](#), [7](#), [14](#)
- [65] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [3](#)
- [66] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision (IJCV)*, 2022. [3](#)
- [67] Yang Zhou, Xintong Han, Eli Shechtman, Jose Echevarria, Evangelos Kalogerakis, and Dingzeyu Li. Makeltalk: speaker-aware talking-head animation. *ACM Transactions on Graphics (TOG)*, 39(6):1–15, 2020. [1](#), [2](#), [5](#), [6](#), [7](#), [13](#), [14](#), [15](#), [16](#)
- [68] Feng Zhu, Zongxin Yang, Xin Yu, Yi Yang, and Yunchao Wei. Instance as identity: A generic online paradigm for video instance segmentation. In *European Conference on Computer Vision*, pages 524–540. Springer, 2022. [2](#)

Supplementary Material




Figure 9. Additional emotional expressions generated by EAT. EAT produces realistic and diverse facial expressions with corresponding emotional guidance. Please zoom in for a better view. Source images are from CREMA-DI[6] and MakeItTalk[67].

A. The Networks Details

We provide additional details of our network architecture and training procedure. It should be noted that the Keypoint Detector (D_k) and RePos-Net networks are primarily derived from OSFV [56]. For more information, interested readers may refer to OSFV [56].

Audio-to-Expression Transformer. We use the Audio-to-Expression Transformer (A2ET) to transfer the audio to 3D latent expression deformation sequences. The A2ET consists of an encoder and a decoder, both with 6 transformer layers and 8 heads. The feed-forward layer has a dimension of 1024. Each token is a 128-dim vector. The expression deformation vector (E_i) is predicted by the feature of the central frame i . However, directly optimizing the 3D expression motions leads to convergence problems in network training. To address this issue and bridge the gap between the 3D expression deformation and the audio features, we use principal component analysis (PCA) to reduce the dimensionality of E_i from 45 to 32. Specifically, we calculate the matrix of principal eigenvalues U and mean vector M from the training set. Then the expression deformation vector is obtained by projecting the predicted PCA using the following equation:

$$E_i = PE_i * U^T + M, \quad (8)$$

where PE_i is the predicted PCA and E_i is the expression deformation, which is used to modify the neutral 3D keypoints to generate the expressive face. As the number of keypoints is 15, the shape of E_i is (15, 3).

Emotion Mapper. We propose an emotion mapper that produces emotional tokens to guide the generation of emotional expressions. As shown in Fig. 10(a), the emotion mapper M consists of several shared and unshared multi-layer perceptrons (MLP) layers. It takes a 16-dim latent

code z as input and outputs seven emotional tokens e_0, e_1, \dots, e_6 . The first token e_0 serves as the emotional guidance for the emotional adaptation module (EAM), which modifies the features of the audio-to-expression transformer (A2ET). The remaining six tokens e_1, \dots, e_6 are fed to the corresponding transformer layer of A2ET as deep emotional prompts. The Emotional Deformation Network (EDN) then uses all these tokens and the latent source representation to generate the emotional deformation ΔE .

Emotional Deformation Network. The Emotional Deformation Network (EDN) learns the emotional deformation ΔE using the same architecture as the A2ET encoder, which has six transformer layers. Fig. 10(b) shows the input and output of EDN. It takes the latent source representation d and the emotional guidance tokens e_0, e_1, \dots, e_6 as input, and extracts their features $f_d, f_{e_0}, \dots, f_{e_6}$. Then it applies global average pooling to the emotion-related features f_{e_0}, \dots, f_{e_6} and uses an MLP layer to obtain the final emotional deformation ΔE .

Emotional Adaptation Module. The emotional adaptation module (EAM) consists of two multi-layer perceptrons (MLPs). As shown in Fig. 10(c), given the input feature x and the emotional token e_0 , we extract the weight vector γ and the bias vector β using MLPs. They have the same dimension as the input x . With the channel-wise multiplication operation F_s and channel-wise addition, we obtain the output x' .

Parameter Efficiency Analysis. Our Deep Emotional Prompts, EDN and EAM only require about 7% of the parameters compared to the whole network. The emotion mapper, which generates deep emotional prompts for eight emotions, has most of the parameters. In addition, EDN and EAM consume less than 2%. These parameters are 13.9M. This is half of the emotional network of EAMM [22], which has 27.9M parameters.

B. Training and Testing Details

Training Details. We use the MEAD dataset and 8k emotional video clips from Voxceleb2 [11] with various facial expressions to learn the enhanced latent keypoints. We also use roughly 21k emotional images from AffectNet [40] to train emotional expression generation. Due to the lack of corresponding neutral faces, we generate neutral faces paired with emotional images by using Ganimation [44]. We train our EAT with Adam [28] with $\beta_1 = 0.5$ and $\beta_2 = 0.999$. The learning rate is set to 1.5×10^{-4} for A2ET and 2×10^{-4} for other modules. In the first stage, we train A2ET with only latent loss first to obtain a good initialization, and then we train it with full loss. To improve gener-

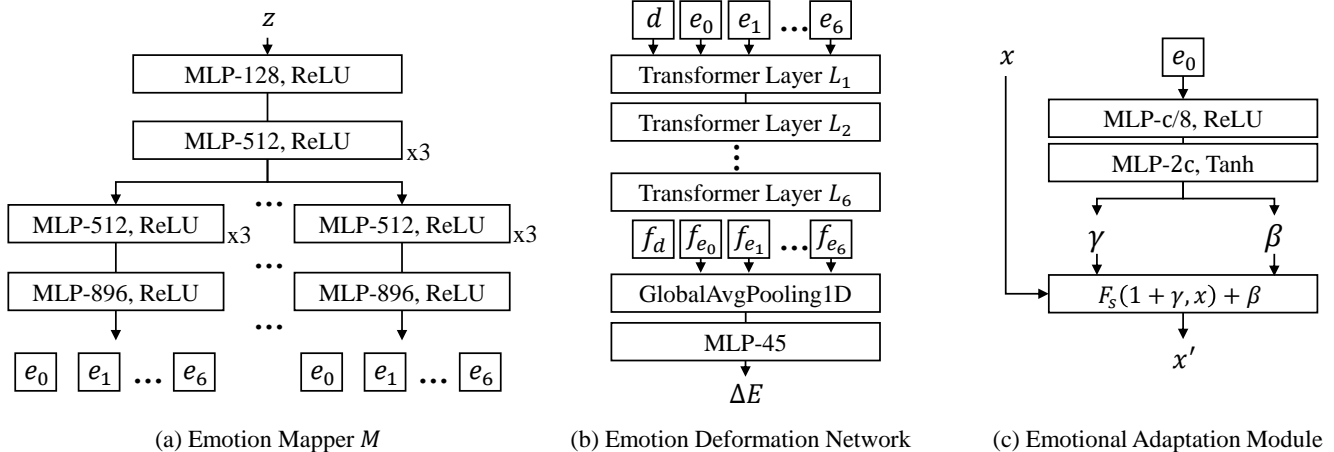


Figure 10. More network architectures of our EAT model.

	Happy	Angry	Disgusted	Fear	Sad	Neutral	Surprised	Contempt	Average
Wav2Lip [43]	0.00	25.64	0.00	0.00	0.00	91.25	0.00	0.00	17.87
MakeItTalk [67]	0.00	25.64	0.00	0.00	0.00	75.00	0.00	0.00	15.23
AVCT [55]	0.83	25.64	0.00	0.00	0.00	69.38	0.00	10.08	15.64
EAMM [22]	23.33	84.48	9.40	0.00	0.00	98.13	94.02	72.27	49.85
Pretrain (Ours)	35.00	11.97	0.00	0.00	49.17	38.75	0.00	59.66	25.18
EAT (Ours)	84.17	100.00	48.72	16.52	49.17	100.00	100.00	94.96	75.43

Table 7. Quantitative evaluation of the emotion classification in the MEAD dataset.

Weight	PSNR \uparrow	M/F-LMD \downarrow	Sync \uparrow	Acc $_{emo}\uparrow$
w/o	21.49	2.27/2.46	8.02	76
EAT	21.79	2.22/2.43	8.22	67

Table 8. Ablation study of EDN weight initialization. The weight initialization of EDN with the A2ET encoder promotes the performance of EAT.

alization, we use the Voxceleb2 and MEAD datasets, which contain about 225k video clips. In the second stage, we finetune efficient adaptation modules with only the MEAD dataset, which has about 10k video clips. We test our model on LRW [12] and MEAD [54] dataset.

Testing Details and Protocol. When testing LRW, the input is the first frame, and the transformation starts from the first frame. Therefore, the relative offsets of the latent keypoints are used. When testing MEAD, due to the variation in facial expressions, which is unrelated to the neutral source image, the predicted latent keypoints are used.

To ensure accurate evaluations, we crop and align [7] the faces before calculating these metrics: PSNR, SSIM, FID, M-LMD, and F-LMD. As for synchronization confidence, we preprocess the generated videos with reference to PC-AVS [64].

C. Additional Experimental Results

Additional baseline results. As shown in Figure 11, we compare our EAT results with several baseline methods. Our results are more pleasant than those of MakeItTalk [67] and Wav2Lip [43], which do not consider emotional expression in talking heads. Additionally, our EAT achieves emotion control compared to the pretrained A2ET network. Videos are included in the supplementary material for reference.

Various emotional expressions To validate the diversity of emotional expressions generated by EAT, we present six different emotional results in Fig. 9. Compared to *Neutral* emotion, emotional expressions result in different modifications to facial elements, such as lip corners, eyes, and brows. We present the quantitative results of emotion classification in Table 7. We notice that EAT works significantly better on *Happy*, *Sad*, *Disgusted*, and *Contempt* than other methods. This is because our method can capture mouth details and these emotions can be more clearly reflected by the lips. As for *Neutral*, *Angry*, and *Surprised*, EAMM [22] performs well since these emotions are more apparent on the eyes and brows. And EAT can also achieve better performance in these emotions. However, all methods perform poorly on *Fear* emotion. It may be because *Fear* and *Surprise* are similar and diffi-

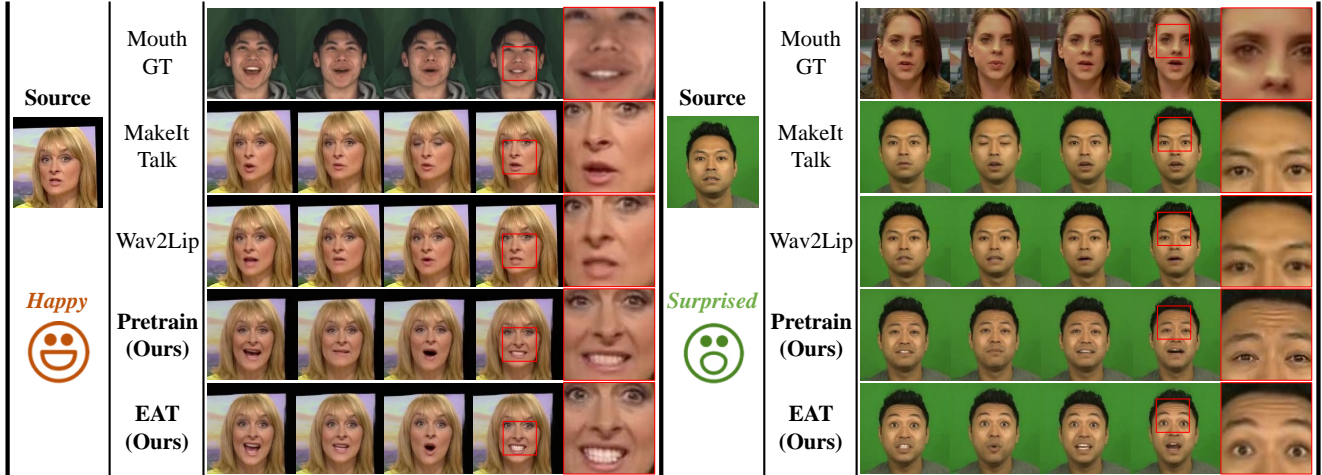


Figure 11. More Qualitative results. We compare with more baselines, such as MakeItTalk [67], Wav2Lip [43], and our pretrained model.

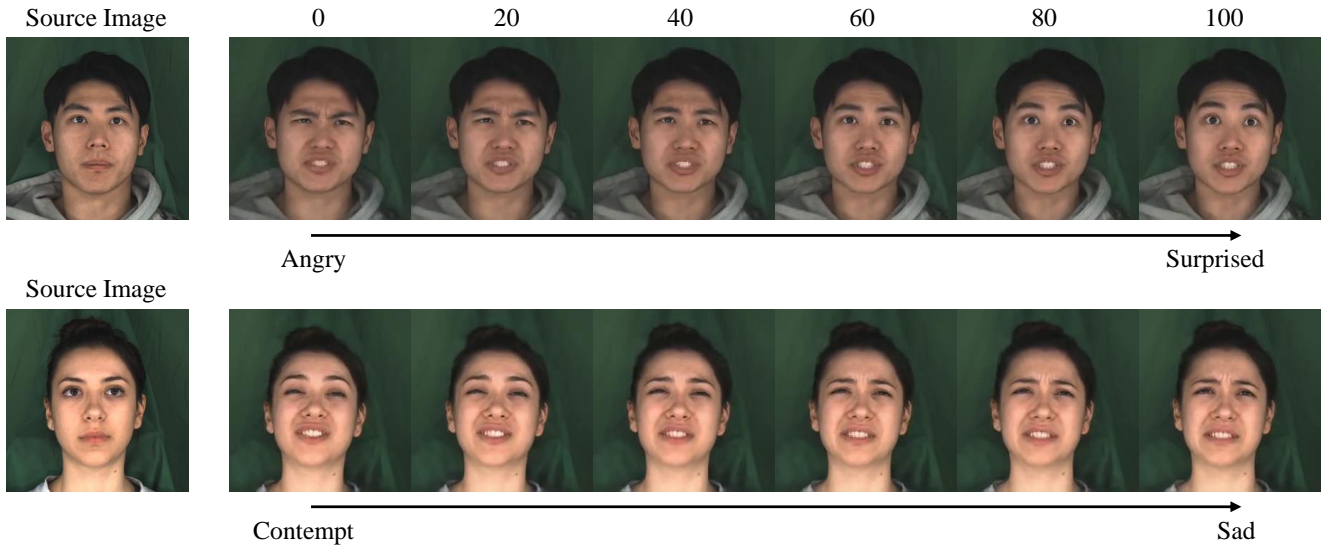


Figure 12. Emotion interpolation. The top row is the emotion interpolation results between *Angry* and *Surprised*. The bottom row is the results between *Contempt* and *Sad*. The neutral faces are from MEAD [54]

cult to distinguish.

Emotion interpolation. We conduct emotional guidance interpolation on the MEAD test set to verify that the latent space learned by the emotion mapper is continuous, as Fig. 12 shows.

Additional ablation study. We conduct further ablation studies on the weight initialization of EDN. Our results, presented in Table 8, show that using the weight initialization of the A2ET encoder leads to quicker convergence and improved performance in terms of video quality and audio-visual synchronization.

Visual analysis on each component of EAT. To analyze the effect of each component of our model, we show the fear emotion results from (A), (B), and (C), with corresponding accuracy rates of 38.46%, 30.77%, and 15.38% respectively in Fig. 13. Deep emotional prompts help generate intense

emotional expressions that deviate from the Ground Truth (GT). By incorporating EDM and EAM, we achieve greater fidelity toward the GT and higher image quality in terms of PSNR/SSIM. This is attributed to the learning capabilities of EDM and EAM for emotional data. However, it results in reduced emotion intensity and accuracy.

Visualization on the profile faces. To assess the ability of enhanced latent representation in 3D talking-head generation, as shown in Fig. 14, we visualize the talking-head frames generated from the profile faces of MEAD. We test the faces captured from left 30 degrees and right 60 degrees with *Surprised* and *Happy* emotions.

D. Limitations and Future Work.

While EAT is capable of generating emotional talking-head videos with emotional guidance, there are still some




Figure 13. Visualization on each component of EAT.




Figure 14. Visualization on the profile faces of MEAD.




Figure 15. Failure case. The driving audio and poses are from the videos in the first row. The second and third rows display the generated results with *Surprised* emotional guidance. Neutral faces are from MakeItTalk [67] and driving video is from LRW [12].

limitations. Firstly, the diversity of background and head pose in emotional training data can affect the generalizability of our EAT. As shown in Fig. 15, the wrinkles on the forehead are not obvious in these in-the-wild images. This issue could be addressed by more naturalistic and non-acted emotional data [60, 30, 4] and representations with the head prior, such as FLAME[62]. Secondly, effective guidance texts are required to achieve zero-shot generation. This may be due to the limited ability of models trained on image-text

pairs to capture emotional expression, which could affect the performance of zero-shot learning. Thirdly, the eye region, such as eye blinks [53] and gaze direction [16], has not been considered in our work. Finally, the discrete emotion guidance limits the representation ability of our model. It needs to note that facial expressions are not always representative of the internal emotional state [3]. More refined theories of emotion, such as the valence-arousal model, may help generate more realistic emotions. We leave these problems for future work.

E. Ethical Considerations.

Our research is intended for use in virtual human research and entertainment. However, there is a risk that the emotional talking-head generation algorithm could be abused. We strongly recommend that generated talking-head videos be labeled as “fake”. On one hand, our work demonstrates that emotional talking-head generation is technically feasible. On the other hand, fake video detection [14, 38] has attracted significant attention. We would be happy to assist in the development of related research.