

# ExpCLIP: Bridging Text and Facial Expressions via Semantic Alignment

Yicheng Zhong\*, Huawei Wei\*, Peiji Yang\*, Zhisheng Wang

<sup>1</sup> Tencent Technology (Shenzhen) Co.Ltd  
{ajaxzhong, huaweiwei, peijiyang, plorywang}@tencent.com

## Abstract

The objective of stylized speech-driven facial animation is to create animations that encapsulate specific emotional expressions. Existing methods often depend on pre-established emotional labels or facial expression templates, which may limit the necessary flexibility for accurately conveying user intent. In this research, we introduce a technique that enables the control of arbitrary styles by leveraging natural language as emotion prompts. This technique presents benefits in terms of both flexibility and user-friendliness. To realize this objective, we initially construct a Text-Expression Alignment Dataset (TEAD), wherein each facial expression is paired with several prompt-like descriptions. We propose an innovative automatic annotation method, supported by ChatGPT, to expedite the dataset construction, thereby eliminating the substantial expense of manual annotation. Following this, we utilize TEAD to train a CLIP-based model, termed ExpCLIP, which encodes text and facial expressions into semantically aligned style embeddings. The embeddings are subsequently integrated into the facial animation generator to yield expressive and controllable facial animations. Given the limited diversity of facial emotions in existing speech-driven facial animation training data, we further introduce an effective Expression Prompt Augmentation (EPA) mechanism to enable the animation generator to support unprecedented richness in style control. Comprehensive experiments illustrate that our method accomplishes expressive facial animation generation and offers enhanced flexibility in effectively conveying the desired style.

## Introduction

In recent years, speech-driven facial animation has gained importance due to its widespread applications in diverse fields such as gaming, virtual reality, and film production (Zhen et al. 2023). Currently, most research focuses on improving the synchronization between lip movements and speech (Cudeiro et al. 2019; Fan et al. 2021; Chen et al. 2022; Xing et al. 2023). This emphasis only allows for conveying speech content, not style, resulting in a lack of emotional expressions in generated facial animations. A few works (Karras et al. 2017; Daněček et al. 2023) have attempted to integrate emotions into facial animation by providing the model with specific emotional labels or use ref-

erence facial expressions as style guidance. However, they either have limited flexibility in expressing diverse emotions or necessitate searching for a reference image or video, which may be impractical for users.

In this study, we propose to adopt natural language as the style prompt for emotional facial animation generation, which offers both flexibility and user-friendliness. A straightforward approach is to collect animation data with paired text prompts and train a text-guided animation generator. However, the scarcity of such data and the high cost of annotation make this approach infeasible. To address this challenge, we propose a novel CLIP-based model called ExpCLIP in this paper. ExpCLIP is designed to learn an embedding space where the representations of text and facial expressions are semantically aligned. By leveraging the capabilities of ExpCLIP, we can train the animation generator by specifying representative facial expressions as prompts, which can be easily extracted from animations, and utilize text prompts for inference purposes.

To train ExpCLIP, a large-scale text-expression dataset is required. However, currently available datasets (Wang et al. 2020; Kollias 2022) only have limited tag-level emotion labels. To address this issue, we propose a novel automated annotation method to construct a Text-Expression Aligned Dataset (TEAD). Specifically, we leverage the visual understanding capability of ChatGPT (OpenAI 2023) to accomplish the annotation task. We find that LLMs are capable of describing the facial expressions corresponding to an emotional text. By harnessing the power of ChatGPT, we collect a rich emotional corpus and use meticulously engineered prompts to ask the ChatGPT to output the corresponding description of facial expressions. Here, we use activated facial Action Units (AUs) to describe facial expressions.

Building upon ExpCLIP, which is trained on TEAD, we propose an emotion-controllable facial animation generator. During the training phase of the generator, we employ a self-attention module to extract the expression prompt from an animation clip. Subsequently, the expression prompt is fed into ExpCLIP to obtain the emotion embedding, which is then fused with the input speech to generate the target facial animation. During the inference stage, we can use text prompts to achieve the desired style control, which is attributed to the alignment of facial expressions and text embeddings achieved by ExpCLIP. We show several exam-

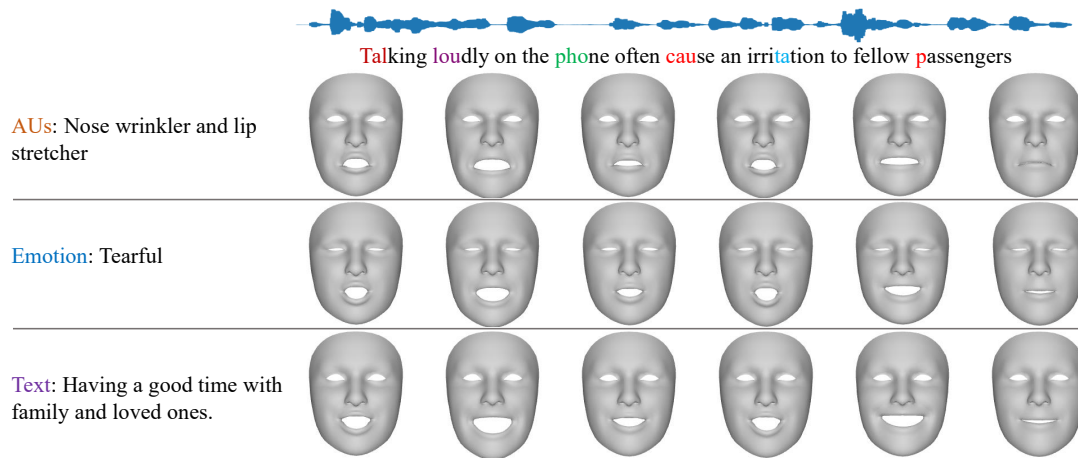


Figure 1: Illustration of our text-guided emotional speech-driven facial animation generation, our approach accommodates diverse textual inputs for style control, encompassing AUs, emotion tags, and other forms of natural language.

ples in Figure 1. Notably, our framework can be readily extended to support facial images as input prompts. This can be achieved by augmenting ExpCLIP with an image encoder and training it on paired facial images and facial expressions. The paired data can be easily obtained with monocular 3D face reconstruction methods (Guo et al. 2022; Lei et al. 2023; Chai et al. 2023).

Moreover, we propose an effective Expression Prompt Augmentation(EPA) mechanism to enable the animation generator to handle unseen emotions. This mechanism involves incorporating random perturbations to the expression prompt and devising a lip motion constraint to ensure the generated lip motions remain consistent with the original motions. The underlying assumption is that when individuals articulate the same sentence with different emotions, their lip motions tend to be consistent.

In summary, the main contributions of our research are:

- We leverage the visual understanding capability of ChatGPT to propose an automatic annotation method for inferring facial expressions from emotional text. This enables us to construct a large-scale text-expression aligned dataset.
- We propose ExpCLIP, which is capable of aligning the semantic representations of text and facial expressions, empowering the inference of emotional styles from natural language descriptions.
- We present the first attempt to use natural language text as prompts to achieve flexible and controllable emotional speech-driven facial animation generation.

## Related Work

### Speech-Driven Face Animation

This field has witnessed a surge of substantial efforts in recent years. VOCA (Cudeiro et al. 2019) interprets the mapping from speech to animation as a regression problem. FaceFormer (Fan et al. 2021) leverages transformers to capture the long-term dependencies inherent in

speech. Meshtalk (Richard et al. 2021) prioritizes addressing the model’s scalability and realism by segregating audio-correlated and audio-uncorrelated information. However, all these methods primarily focus on enhancing lip synchronization, falling short in conveying emotional expressions.

### Emotion Guided Generation

The incorporation of emotional expressions has been considered in recent works. (Karras et al. 2017) uses a trainable variable to represent the emotion state and (Daněček et al. 2023) uses predefined emotion labels to guide the animation generation. Similar approaches have also been observed in several works (Sadoughi and Busso 2019; Ji et al. 2021; Wu et al. 2021; Liang et al. 2022; Sinha et al. 2022) related to talking face generation. (Ji et al. 2022) and (Ma et al. 2023) propose to extract emotional information from reference videos. (Wang et al. 2023) utilize a static facial image as an emotional condition for styled talking face generation. These methods either have limited flexibility in expressing diverse emotions, or necessitate searching for a reference image or video, which may be impractical for users.

### CLIP-Based Content Synthesis

CLIP (Radford et al. 2021) has demonstrated its efficacy text guided image editing (Rombach et al. 2022; Ramesh et al. 2023; Schaldenbrand, Liu, and Oh 2021). It can also be extended to the integration of text with other modalities, such as employing CLIP for 3D motion generation (Tevet et al. 2022). However, to the best of our knowledge, no prior work has utilized CLIP for 3D facial animation generation. In this paper, we propose, for the first time, the use of CLIP for text-guided speech-driven facial animation generation.

## Method

In this section, we first present the construction of TEAD, followed by an overview of the training process for ExpCLIP. Finally, we introduce the proposed text-guided speech-driven facial animation method.

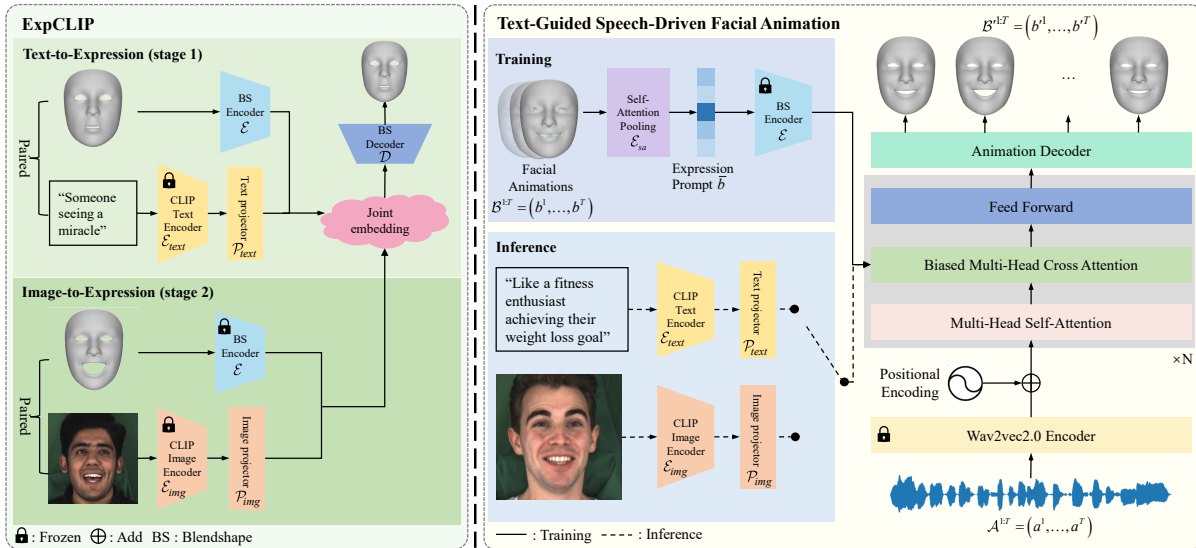


Figure 2: Overview of our framework. We first train ExpCLIP to establish semantic alignment between text ,facial expressions and facial images. Then, we employ it for the generation of emotional speech-driven facial animation. The animation generator supports both images and natural language prompts for emotion control.

### Text-Expression Aligned Dataset

In order to train a model capable of aligning the semantic representations of natural language and facial expressions, a text-expression dataset is required. However, currently available datasets only provide limited emotion labels (Cao et al. 2014; Wang et al. 2020), which are insufficient for achieving the desired alignment at a fine-grained level. To address this limitation, we propose a Text-Expression Aligned Dataset (TEAD), which is automatically constructed with the assistance of ChatGPT(OpenAI 2023).

**Motivation** To generate paired emotional text and corresponding facial expression data using ChatGPT, it is imperative to represent facial expressions as textual descriptions. Fortunately, the Facial Action Coding System (FACS) (Ekman and Friesen 1978) offers a systematic approach to describe human facial movements by decomposing facial expressions into independent AUs, with each AU possessing an exhaustive textual description. Therefore, our objective is to let ChatGPT generate corresponding activated AUs from the emotional text. This necessitates that ChatGPT demonstrate cross-modal understanding capability, that is, the ability to “**imagine**” corresponding facial expressions that depict the emotion in the text, and subsequently translate the expressions into activated AUs. Through our testing, we find that ChatGPT can successfully accomplish this task with a carefully designed prompt through in-context learning.

**Automatic Data Generation** We utilize the abundant corpus from text emotion classification tasks (Mohammad and Bravo-Marquez 2017), which encompass rich real-world human emotions. The emotional transcripts are fed into the ChatGPT for text emotion classification and detection of activated AUs. Specifically, given an emotional transcript  $t$ , the ChatGPT predicts its emotional tags  $e$ , where  $e$  contains several emotional labels, usually 3 to 5, and the activated AUs




Figure 3: An example of data generation of TEAD.

with a one-hot vector  $u = \{u_i, i = 1, \dots, N_u\}$ .  $N_u$  is 36 in our work. Note the combination of these AUs enables a wider range of emotions compared to existing datasets that have only a few emotion labels. We further prompt the ChatGPT to describe situations that may evoke inferred emotions, generating sentence-level labels  $s$ . An instance is shown in Figure 3.

In addition, we enlist a professional facial animator to devise a mapping rule to convert AUs into blendshape weights. Then, we transform the AUs vector  $u$  into a set of blendshape weights  $b \in \mathbb{R}^{52}$ , enabling us to leverage publicly available datasets that represent facial expressions as blendshape weights. Consequently, TEAD can be represented as a set of quadruples:  $\mathcal{T} = \{(t, e, b, s) | i = 1, \dots, N_T\}$ . About 50,000 quadruples are included.

### ExpCLIP

We utilize TEAD to train a CLIP-based model, named ExpCLIP. As illustrated in Figure 2, ExpCLIP is an autoencoder framework designed to align multimodal signals, encompassing diverse human facial expressions, textual descriptions of emotions, and realistic facial images.

**Text-to-Expression** We propose a blendshape encoder  $\mathcal{E}$  that maps a given set of blendshape weights  $b$  into an embedding. Subsequently, a decoder  $\mathcal{D}$  reconstructs the blendshape weights from the embedding. In parallel, we utilize a CLIP text encoder  $\mathcal{E}_{text}$  based on the work of (Radford et al. 2021), along with a text projector  $\mathcal{P}_{text}$  to map the emotion text, sampled from  $\{t, e, s\}$ , into the joint embedding space. Due to the strong generalization capability of pre-trained CLIP, we keep  $\mathcal{E}_{text}$  frozen during training to leverage its acquired extensive knowledge.

**Image-to-Expression** We propose an extension to incorporate facial images as emotion prompts by integrating an image encoder into ExpCLIP and aligning the embeddings of images and expressions. A SOTA 3D face reconstruction method (Guo et al. 2022) is used to create paired blendshape weights for facial images. We leverage the pre-trained image encoder  $\mathcal{E}_{img}$  from CLIP and employ an image projector  $\mathcal{P}_{img}$  to map facial images into the joint embedding space. During this step, we only finetune  $\mathcal{P}_{img}$ , while keeping the the weights of  $\mathcal{E}_{img}$ ,  $\mathcal{E}$ , and  $\mathcal{D}$  fixed.

**Objective Function** ExpCLIP is trained via three types of losses: auto-encoder reconstruction loss, embedding alignment loss, and cross-modal reconstruction loss:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{ae} + \lambda_2 \mathcal{L}_{emb}^d + \lambda_3 \mathcal{L}_{cross}^d, \quad (1)$$

where  $d \in \{text, image\}$ . Note that  $d = text$  indicates the text-to-expression training, and  $d = image$  refers to the image-to-expression finetuning. The auto-encoder reconstruction loss  $\mathcal{L}_{ae}$  measures the L2 distance between the input blendshape weights  $b$  and the predicted ones:

$$\mathcal{L}_{ae} = \mathbb{E}_{b \sim \mathcal{T}} \|\mathcal{D}(\mathcal{E}(b)) - b\|. \quad (2)$$

We align the embeddings of text (or images) and expressions using cosine embedding loss:

$$\mathcal{L}_{emb}^d = 1 - \cos(\mathcal{P}_d(\mathcal{E}_d(S)) - \mathcal{E}(b)), \quad (3)$$

where  $S$  represents text in  $\{t, e, s\}$  when  $d = text$ , and facial images when  $d = image$ .

To better improve the multi-modal alignment, we propose a cross-modal reconstruction loss, which enforces the blendshape weights reconstructed using text (or image) embeddings close to its paired ones, namely:

$$\mathcal{L}_{cross}^d = \mathbb{E}_{b \sim \mathcal{T}} \|\mathcal{D}(\mathcal{P}_d(\mathcal{E}_d(S))) - b\|. \quad (4)$$

We propose several strategies to enhance the model’s generality. They include fully exploiting the TEAD dataset by randomly extracting samples from  $t$ ,  $e$ , and  $s$  as text inputs for training. Additionally, we borrow text augmentation techniques from the natural language processing literature (Wei and Zou 2019), such as stop-word removal, synonym replacement, and sentence shuffling. Furthermore, we augment the blendshape weights by applying minor random perturbations. Experiments show all these strategies improve the model’s robustness.

### Text-Guided Speech-Driven Facial Animation

Our aim is to automatically generate emotionally expressive facial animation, where the content is determined by speech and the emotional style is controlled by a text prompt. The overview of our method is shown in Figure 2.

**Training Workflow** Leveraging the semantic alignment between text and facial expressions achieved by ExpCLIP, we can employ expression prompts for training and text prompts for inference. Yet, annotating expression prompts for each animation clip is costly. To address this, we propose a self-attention mechanism to automatically extract representative expressions from the animation clips. Consequently, the training workflow of our model is as follows:

Let  $\mathcal{A}^{1:T} = (a^1, \dots, a^T)$  denotes a sequence of speech snippets.  $\mathcal{B}^{1:T} = (b^1, \dots, b^T)$  is the synchronized facial animation, each frame is represented by a set of blendshape weights  $b^t$ . We propose a transformer-based self-attention pooling module, denoted as  $\mathcal{E}_{sa}$ , to derive the attention weights for individual frames within the animation clip. Subsequently, we aggregate all frames of the clip according to the attention weights, resulting in the generation of the expression prompt  $\bar{b}$ . Next, we feed the expression prompt  $\bar{b}$  into the expression encoder  $\mathcal{E}$  of ExpCLIP to obtain the style embedding. During the training process, the parameters of  $\mathcal{E}$  are kept frozen. Simultaneously, we utilize a pre-trained wav2vec2.0 model (Baevski et al. 2020) to convert raw waveform input into contextualized speech features. We employ a transformer decoder to predict  $\mathcal{B}'^{1:T} = (b'^1, \dots, b'^T)$  from speech features, with style embedding incorporated into the decoding process via cross-attention.

A simple L1 loss is utilized for reconstruction:

$$\mathcal{L}_{rec} = \frac{1}{T} \sum_{t=0}^T |b'^t - b^t| \quad (5)$$

**Expression Prompt Augmentation** The currently publicly accessible speech-driven facial animation datasets exhibit insufficiency in terms of emotional richness. They are typically constrained to a limited number of coarse-grained emotion labels, thereby impeding the model’s ability to handle unseen fine-grained emotions. To address this issue, we propose an Expression Prompt Augmentation (EPA) mechanism. Specifically, we add perturbations to the expression prompts. In order to make the perturbed prompts show certain emotions instead of random weird expressions, the perturbations are obtained by randomly sampling a facial expression  $b_{aug}$  from TEAD. Subsequently, we blend  $b_{aug}$  into the original expression prompt using a random weight  $\lambda$ .

$$\bar{b}_{aug} = (1 - \lambda)\bar{b} + \lambda b_{aug}, \lambda \in [0, 1] \quad (6)$$

Due to the lack of corresponding animations for the perturbed expression prompts, we devise two loss functions to ensure that the generated animations exhibit accurate lip movements and emotional styles. We posit that when a person utters the same sentence with different emotions, their lip movements remain fundamentally consistent, meaning that the displacement between adjacent frames remains consistent. Leveraging this assumption, we formulate a lip motion loss to capture and enforce this consistency in the generated animations. We represent the animation for the perturbed prompt as  $\mathcal{B}'_{aug} = (b'_{aug}^1, \dots, b'_{aug}^T)$ , and the lip

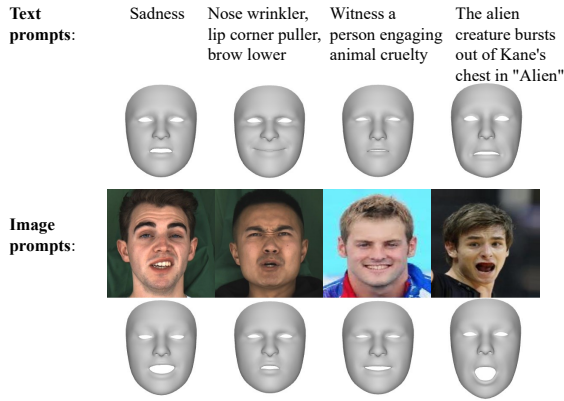


Figure 4: Multi-modal alignment. Row 1&2: text-to-expression. Row 3&4: image-to-expression. Note the last two columns are *out-of-domain* samples.

motion loss is defined as:

$$\mathcal{L}_{lm} = \frac{1}{T-1} \sum_{t=0}^{T-1} |(b'^{t+1}_{aug} - b^t_{aug}) - (b^{t+1} - b^t)| \quad (7)$$

$\mathcal{L}_{lm}$  only ensures the correctness of lip motion in animations generated based on perturbed prompts, but it does not guarantee the desired style. To address this limitation, we propose a style loss:

$$\mathcal{L}_{style} = \|\mathcal{E}(\mathcal{E}_{sa}(\mathcal{B}'_{aug})) - \mathcal{E}(\bar{b}_{aug})\| \quad (8)$$

$\mathcal{L}_{style}$  ensures the consistency between the emotion of the generated animation  $\mathcal{B}'_{aug}$  and the augmented expression prompt  $\bar{b}_{aug}$ .

**Inference Phase** Since the embeddings of text and facial expressions are semantically aligned by ExpCLIP, we can employ text descriptions as prompts to control the emotions of the generated animations. There is no strict requirement for the text to adhere to a specific format. You have the freedom to express your text prompts in any manner you prefer. For instance, you may provide a precise description of the desired facial expression or simply describe the mood you want to express. Furthermore, in cases where verbal text may not effectively convey the intended emotions, the option of employing a reference facial image as a prompt is also available. This capability is facilitated by ExpCLIP’s alignment of facial images and facial expressions.

## Experiments

In this section, we first introduce the used datasets and implementation details. Subsequently, we present the capabilities of ExpCLIP in multimodal alignment and discuss several key factors for training ExpCLIP. Finally, we show the promising results of style-controllable speech-driven facial animation, accompanied by detailed comparative experiments and ablation studies.

### Datasets

**TEAD** We train ExpCLIP using the proposed TEAD, which consists of 50,000 quadruples. Each quadruple includes text, a set of emotion tags, AUs, blendshape weights,

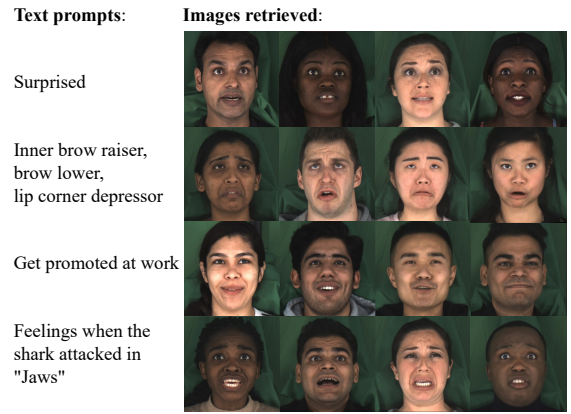


Figure 5: Text-to-image retrieval on MEAD-3D.




Figure 6: (a) t-SNE of expression embeddings in TEAD. (b) smooth interpolations between two distant expressions.

and situation sentences. We use 90% of the data for training and the remaining 10% for testing the text-expression alignment of ExpCLIP.

**MEAD-3D** To support the image-expression alignment of ExpCLIP, we generate image-expression paired data based on MEAD. MEAD is a talking-face video corpus featuring 60 actors talking with 8 different emotions at 3 different intensity levels. We sample 150,000 images from MEAD and use a SOTA 3D monocular face reconstruction method (Guo et al. 2022) to obtain the 3D facial meshes. Then we compute the blendshape weights corresponding to each mesh using the SLSQP solver from SciPy (Virtanen et al. 2020). We use 90% of the data for training and the remaining 10% for testing the image-expression alignment effect of ExpCLIP.

**BEAT** We use BEAT (Liu et al. 2022) to train the speech-driven facial animation generator. BEAT comprises 76 hours of speech data, paired with 52D facial blendshape weights. The dataset is collected from 30 speakers, who perform in 8 distinct emotional styles and across 4 different languages. For our experiments, we exclusively employ speech data from English speakers, which totals approximately 35 hours.

### Implementation Details

Our framework is implemented by Pytorch (Paszke et al. 2019). For ExpCLIP, we train a transformer auto-encoder (Vaswani et al. 2017) with 8 layers for both the encoder  $\mathcal{E}$  and  $\mathcal{D}$ . The text encoder and image encoder from CLIP-ViT-B/32 are utilized. We set the values of  $\lambda_1 = 1, \lambda_2 = \lambda_3 = 10$ . ExpCLIP is trained with a learning rate of  $1e-5$  and a

Datasets	TEAD	BEAT
w/o bs aug	0.051	0.069
w/. bs aug	<b>0.012</b>	<b>0.021</b>

Table 1: MSE of blendshape weights reconstruction on the test set of TEAD and BEAT.

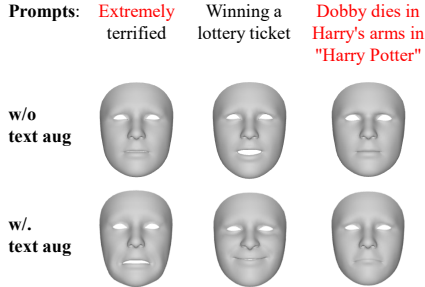


Figure 7: Qualitative results of ablation study for text augmentation. Red words indicate *out-of-domain* texts.

batch size of 256. For the animation generator, we employ a 4-layer transformer encoder as the self-attention pooling module. An 8-layer transformer decoder is used to modulate the speech features and style embeddings. The animation decoder consists of 2-layer fully connected layers. Each training sample has a duration of 64 frames with FPS=15. The entire framework is trained using the Adam optimizer (Kingma and Ba 2014) on a single A100 GPU.

## ExpCLIP

**Multi-modal Alignment** ExpCLIP aligns expressions, text, and images into a joint embedding space. We examine this capability through three types of tasks: text-to-expression, text-to-image, and image-to-expression.

Text-to-expression transforms text to blendshape weights, which are rendered to meshes for better visualization. In-domain text from the test set and out-of-domain text are collected for evaluation. Image-to-expression can be treated as a 3D face reconstruction task, which transforms images to blendshape weights. Figure 4 shows that ExpCLIP exhibits the ability to generate subtle and nuanced expressions, accommodating various types of textual input such as emotion tags, AU descriptions, and sentences. Furthermore, ExpCLIP also excels in the task of recovering intricate facial expressions from in-the-wild images.

As for text-to-image, we employ a text-based image retrieval task to evaluate its performance. We extract image embeddings of the test set of MEAD-3D and employ cosine similarity to retrieve the images that are semantically closest to the embedding of the given text. Figure 5 indicates ExpCLIP achieves a remarkable alignment between textual and visual semantics.

**Expression Manifold Smoothness** We demonstrate the smoothness of the learned expression manifold in Figure 6. We obtain embeddings of facial expressions in the test

set of TEAD. These embeddings are then projected onto a 2D space using t-SNE (Van der Maaten and Hinton 2008). Subsequently, we sample pairs of distant points and interpolate between them. The interpolated values are passed through the decoder of ExpCLIP to reconstruct the corresponding blendshape weights. As observed, it illustrates ExpCLIP achieves smooth semantic transition between distinct facial expressions.

**Ablation Study** We conduct an ablation study to validate the effectiveness of blendshape augmentation. The mean squared error (MSE) of blendshape weight reconstruction on the test set of TEAD and BEAT is presented in Table 1. The results demonstrate that blendshape augmentation significantly contributes to reducing the reconstruction error of ExpCLIP. To evaluate the influence of text augmentation, we train a model without text augmentation on TEAD. The qualitative outcomes are depicted in Figure 7. It is evident that text augmentation enhances the consistency between the generated expressions and out-of-domain text prompts.

## Emotional Speech-Driven Facial Animation

**Baselines** We compare our method with SOTA emotion-controllable talking face generation methods including: **EAMM** (Ji et al. 2022), **StyleTalk** (Ma et al. 2023) and **PD-FGC** (Wang et al. 2023). We employ the 3D face reconstruction method (Guo et al. 2022) to reconstruct facial meshes from their generated talking face videos for comparative analysis. All these methods employ facial images or video templates as emotional prompts. For comparison purposes, we manually annotate textual descriptions for each template. Notably, as both PD-FGC and our method accept image-based emotional prompts, we can utilize this setting for comparison.

Due to the lack of appropriate quantitative metrics to characterize the precision of emotion control in animation generation, we only conduct qualitative evaluations and user studies to evaluate the above methods.

**Qualitative Results** The qualitative results are illustrated in Figure 8. The image and video templates are extracted from the test set of MEAD. Both BEAT and MEAD provided speech samples for evaluation.

As can be seen, our method ensures a high level of consistency between the emotional expression in the generated animation and the provided prompt, while also achieving accurate lip synchronization. Even in the more challenging setting of using text as an emotional prompt, our approach achieves high accuracy in emotion control. In comparison, EAMM exhibits inferior lip synchronization and demonstrates significant inconsistency between the emotion of the generated animation and the reference video. Although StyleTalk and PD-FGC achieve satisfactory lip synchronization, their accuracy in emotion control falls short compared to our approach. Additionally, their control methods are inconvenient for users as they require searching for desired emotional reference templates. In contrast, our text-based control approach is flexible and easy to manipulate.

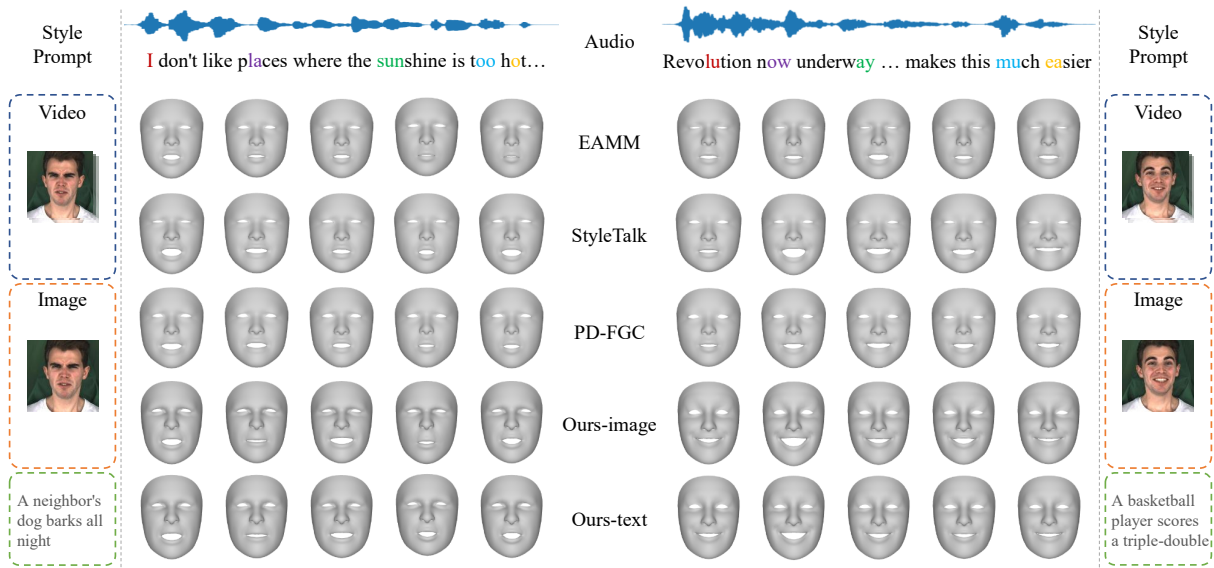


Figure 8: Qualitative comparisons with SOTA stylized talking face generation methods. Note that the speaking style of our method can be guided by text description or emotional image.

Methods	EAMM	StyleTalk	PD-FGC	Ours-image	Ours-text
Lip sync	2.01	3.55	3.49	3.59	<b>3.63</b>
$EC_v$	2.11	3.63	3.55	<b>3.72</b>	3.65
$EC_t$	1.99	3.48	3.37	3.75	<b>3.93</b>

Table 2: Results of the user study

**User Study** We further conduct user studies to evaluate the performance of the comparative methods. We create 10 emotional animations for each method, accompanied by corresponding reference videos or images, as well as text prompts. We invite 20 volunteers to rate these methods from 1 to 5, with higher scores indicating better performance. The volunteers are asked to rate the methods based on the following three aspects: 1. Lip synchronization, 2. Emotion consistency with the reference video/image, and 3. Emotion consistency with the text prompt.

As shown in Table 2, our performance in lip synchronization exceeds that of StyleTalk and PD-FGC, and markedly surpasses EAMM. Regarding emotion consistency with video/image ( $EC_v$ ), our text-based control method approximates the level of StyleTalk and PD-FGC, while our image-based control method exhibits superior performance compared to other methods. As for emotion consistency with text ( $EC_t$ ), our text-based control method outperforms all other methods, and our image-based control method also surpasses other video or image-based approaches. The above results demonstrate that our method not only achieves precise lip synchronization but also enables accurate emotion control.

**Ablation Study** We conduct ablation studies to validate the effect of EPA and the proposed style loss  $\mathcal{L}_{style}$ , the qualitative results are shown in Figure 9. Notably, to better visualize the differences between the comparisons, we




Figure 9: Ablation studies of EPA and style loss.

utilize blendshape weights as the prompt, rather than text. As illustrated, when there is no EPA, if an unseen emotion is input, the resulting animation exhibits minimal expression. This is due to the limited generalization ability of the model toward unfamiliar emotions. However, the integration of EPA markedly enhances the emotional expressiveness of the animation. Despite this improvement, the consistency with the expression prompt remains somewhat less than ideal. The incorporation of style loss  $\mathcal{L}_{style}$  further intensifies the emotional impact.

## Conclusion

This paper introduces, for the first time, text-guided emotional speech-driven facial animation. To achieve this, a large-scale text-expression dataset TEAD is proposed, and ExpCLIP is trained on this dataset to align features of text and expressions. Experimental results demonstrate that the proposed framework achieves high accuracy and flexibility in emotional-controlled animation generation.

## References

- Baevski, A.; Zhou, Y.; Mohamed, A.; and Auli, M. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems*, 33: 12449–12460.
- Cao, H.; Cooper, D. G.; Keutmann, M. K.; Gur, R. C.; Nenkova, A.; and Verma, R. 2014. CREMA-D: Crowd-Sourced Emotional Multimodal Actors Dataset. *IEEE Transactions on Affective Computing*, 5(4): 377–390.
- Chai, Z.; Zhang, T.; He, T.; Tan, X.; Baltrusaitis, T.; Wu, H.; Li, R.; Zhao, S.; Yuan, C.; and Bian, J. 2023. HiFace: High-Fidelity 3D Face Reconstruction by Learning Static and Dynamic Details. *arXiv preprint arXiv:2303.11225*.
- Chen, L.; Wu, Z.; Ling, J.; Li, R.; Tan, X.; and Zhao, S. 2022. Transformer-s2a: Robust and efficient speech-to-animation. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 7247–7251. IEEE.
- Cudeiro, D.; Bolkart, T.; Laidlaw, C.; Ranjan, A.; and Black, M. J. 2019. Capture, Learning, and Synthesis of 3D Speaking Styles. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10093–10103.
- Daněček, R.; Chhatre, K.; Tripathi, S.; Wen, Y.; Black, M. J.; and Bolkart, T. 2023. Emotional Speech-Driven Animation with Content-Emotion Disentanglement. *arXiv preprint arXiv:2306.08990*.
- Ekman, P.; and Friesen, W. V. 1978. Facial action coding system. *Environmental Psychology & Nonverbal Behavior*.
- Fan, Y.; Lin, Z.; Saito, J.; Wang, W.; and Komura, T. 2021. FaceFormer: Speech-Driven 3D Facial Animation with Transformers. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 18749–18758.
- Guo, J.; Yu, J.; Lattas, A.; and Deng, J. 2022. Perspective reconstruction of human faces by joint mesh and landmark regression. In *European Conference on Computer Vision*, 350–365. Springer.
- Ji, X.; Zhou, H.; Wang, K.; Wu, Q.; Wu, W.; Xu, F.; and Cao, X. 2022. Eamm: One-shot emotional talking face via audio-based emotion-aware motion model. In *ACM SIGGRAPH 2022 Conference Proceedings*, 1–10.
- Ji, X.; Zhou, H.; Wang, K.; Wu, W.; Loy, C. C.; Cao, X.; and Xu, F. 2021. Audio-driven emotional video portraits. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14080–14089.
- Karras, T.; Aila, T.; Laine, S.; Herva, A.; and Lehtinen, J. 2017. Audio-driven facial animation by joint end-to-end learning of pose and emotion. *ACM Transactions on Graphics (TOG)*, 36(4): 1–12.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kollias, D. 2022. Abaw: Valence-arousal estimation, expression recognition, action unit detection & multi-task learning challenges. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2328–2336.
- Lei, B.; Ren, J.; Feng, M.; Cui, M.; and Xie, X. 2023. A Hierarchical Representation Network for Accurate and Detailed Face Reconstruction from In-The-Wild Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 394–403.
- Liang, B.; Pan, Y.; Guo, Z.; Zhou, H.; Hong, Z.; Han, X.; Han, J.; Liu, J.; Ding, E.; and Wang, J. 2022. Expressive talking head generation with granular audio-visual control. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3387–3396.
- Liu, H.; Zhu, Z.; Iwamoto, N.; Peng, Y.; Li, Z.; Zhou, Y.; Bozkurt, E.; and Zheng, B. 2022. Beat: A large-scale semantic and emotional multi-modal dataset for conversational gestures synthesis. In *European Conference on Computer Vision*, 612–630. Springer.
- Ma, Y.; Wang, S.; Hu, Z.; Fan, C.; Lv, T.; Ding, Y.; Deng, Z.; and Yu, X. 2023. Styletalk: One-shot talking head generation with controllable speaking styles. *arXiv preprint arXiv:2301.01081*.
- Mohammad, S. M.; and Bravo-Marquez, F. 2017. Emotion Intensities in Tweets. In *Proceedings of the sixth joint conference on lexical and computational semantics (\*Sem)*. Vancouver, Canada.
- OpenAI. 2023. ChatGPT. <https://chat.openai.com/chat>. Accessed: 2023-05-14.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2023. Hierarchical text-conditional image generation with clip latents.(arXiv preprint)(2022). DOI: <https://doi.org/10.48550/ARXIV.2204>.
- Richard, A.; Zollhoefer, M.; Wen, Y.; la Torre, F. D.; and Sheikh, Y. 2021. MeshTalk: 3D Face Animation from Speech using Cross-Modality Disentanglement. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 1153–1162.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Sadoughi, N.; and Busso, C. 2019. Speech-driven expressive talking lips with conditional sequential generative adversarial networks. *IEEE Transactions on Affective Computing*, 12(4): 1031–1044.
- Schaldenbrand, P.; Liu, Z.; and Oh, J. 2021. Styleclipdraw: Coupling content and style in text-to-drawing synthesis. *arXiv preprint arXiv:2111.03133*.

- Sinha, S.; Biswas, S.; Yadav, R.; and Bhowmick, B. 2022. Emotion-controllable generalized talking face generation. *arXiv preprint arXiv:2205.01155*.
- Tevet, G.; Gordon, B.; Hertz, A.; Bermano, A. H.; and Cohen-Or, D. 2022. Motionclip: Exposing human motion generation to clip space. In *European Conference on Computer Vision*, 358–374. Springer.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Virtanen, P.; Gommers, R.; Oliphant, T. E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Burovski, E.; Peterson, P.; Weckesser, W.; Bright, J.; et al. 2020. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature methods*, 17(3): 261–272.
- Wang, D.; Deng, Y.; Yin, Z.; Shum, H.-Y.; and Wang, B. 2023. Progressive Disentangled Representation Learning for Fine-Grained Controllable Talking Head Synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17979–17989.
- Wang, K.; Wu, Q.; Song, L.; Yang, Z.; Wu, W.; Qian, C.; He, R.; Qiao, Y.; and Loy, C. C. 2020. Mead: A large-scale audio-visual dataset for emotional talking-face generation. In *European Conference on Computer Vision*, 700–717. Springer.
- Wei, J.; and Zou, K. 2019. Eda: Easy data augmentation techniques for boosting performance on text classification tasks. *arXiv preprint arXiv:1901.11196*.
- Wu, H.; Jia, J.; Wang, H.; Dou, Y.; Duan, C.; and Deng, Q. 2021. Imitating arbitrary talking style for realistic audio-driven talking face synthesis. In *Proceedings of the 29th ACM International Conference on Multimedia*, 1478–1486.
- Xing, J.; Xia, M.; Zhang, Y.; Cun, X.; Wang, J.; and Wong, T.-T. 2023. Codetalker: Speech-driven 3d facial animation with discrete motion prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12780–12790.
- Zhen, R.; Song, W.; He, Q.; Cao, J.; Shi, L.; and Luo, J. 2023. Human-computer interaction system: A survey of talking-head generation. *Electronics*, 12(1): 218.