

FaceFormer: Speech-Driven 3D Facial Animation with Transformers

Yingruo Fan¹ Zhaojiang Lin^{2†} Jun Saito³ Wenping Wang^{1,4} Taku Komura^{1*}

¹The University of Hong Kong ²The Hong Kong University of Science and Technology
³Adobe Research ⁴Texas A&M University

Abstract

Speech-driven 3D facial animation is challenging due to the complex geometry of human faces and the limited availability of 3D audio-visual data. Prior works typically focus on learning phoneme-level features of short audio windows with limited context, occasionally resulting in inaccurate lip movements. To tackle this limitation, we propose a Transformer-based autoregressive model, FaceFormer, which encodes the long-term audio context and autoregressively predicts a sequence of animated 3D face meshes. To cope with the data scarcity issue, we integrate the self-supervised pre-trained speech representations. Also, we devise two biased attention mechanisms well suited to this specific task, including the biased cross-modal multi-head (MH) attention and the biased causal MH self-attention with a periodic positional encoding strategy. The former effectively aligns the audio-motion modalities, whereas the latter offers abilities to generalize to longer audio sequences. Extensive experiments and a perceptual user study show that our approach outperforms the existing state-of-the-arts. The code and the video are available at: <https://evelynfan.github.io/audio2face/>.

1. Introduction

Speech-driven 3D facial animation has become an increasingly attractive research area in both academia and industry. It is potentially beneficial to a broad range of applications such as virtual reality, film production, games and education. Realistic speech-driven 3D facial animation aims to automatically animate vivid facial expressions of the 3D avatar from an arbitrary speech signal.

We focus on animating the 3D geometry rather than the 2D pixel values, *e.g.* photorealistic talking-head animation [12, 15, 42, 52, 63, 67, 69]. The majority of existing works aim to produce 2D videos of talking heads, given the availability of massive 2D video datasets. However, the




Figure 1. **Concept diagram of FaceFormer.** Given the raw audio input and a neutral 3D face mesh, our proposed end-to-end Transformer-based architecture, dubbed FaceFormer, can autoregressively synthesize a sequence of realistic 3D facial motions with accurate lip movements.

generated 2D videos are not directly applicable to applications like 3D games and VR, which need to animate 3D models in a 3D environment. Several methods [27, 47, 60] harness 2D monocular videos to obtain 3D facial parameters, which might lead to unreliable results. This is because the quality of the synthetic 3D data is bounded by the accuracy of 3D reconstruction techniques, which cannot capture the subtle changes in 3D. In speech-driven 3D facial animation, some 3D mesh-based works [8, 17, 39] formulate the input as short audio windows, which might result in ambiguities in variations of facial expressions. As pointed out by Karras *et al.* [31], the emotional state can be deduced from a longer-term audio context. While MeshTalk [51] has considered a longer audio context by modeling the audio sequence, training the model with Mel spectral audio features fails to synthesize accurate lip motions in data-scarce settings. Collecting lots of pairs of speech and 3D motion capture data is also considerably expensive and time-consuming.

To address the issues about long-term context and lack of 3D audio-visual data, we propose a transformer-based autoregressive model (Fig. 1) which (1) captures longer-term audio context to enable highly realistic animation of the en-

* Corresponding author

† Work done at HKUST

tire face, *i.e.* both upper and lower face expressions, (2) effectively utilizes the self-supervised pre-trained speech representations to handle the data scarcity issue, and (3) considers the history of face motions for producing temporally stable facial animation.

Transformer [58] has achieved remarkable performance in both natural language processing [20, 58] and computer vision [13, 21, 44] tasks. The sequential models like LSTM have a bottleneck that hinders the ability to learn longer-term context effectively [45]. Compared to RNN-based models, transformer can better capture long-range context dependencies based solely on attention mechanisms [58]. Recently, transformer has also made the encouraging progress in body motion synthesis [1, 4, 46] and dance generation [36, 37, 57]. The success of transformer is mainly attributed to its design incorporating the self-attention mechanism, which is effective in modeling both the short- and long-range relations by explicitly attending to all parts of the representation. Speech-driven 3D facial animation has not been explored in this direction.

Direct application of a vanilla transformer architecture to audio sequences does not perform well on the task of speech-driven 3D facial animation, and we thus need to address these issues. First, transformer is data-hungry in nature, requiring sufficiently large datasets for training [32]. Given the limited availability of 3D audio-visual data, we explore the use of the self-supervised pre-trained speech model wav2vec 2.0 [2]. Wav2vec 2.0 has learned rich phoneme information, since it has been trained on a large-scale corpus [43] of unlabeled speech. While the limited 3D audio-visual data might not cover enough phonemes, we expect the pre-trained speech representations can benefit the speech-driven 3D facial animation task in data-scarce settings. Second, the default encoder-decoder attention of transformer can not handle modality alignment, and thus we add an alignment bias for audio-motion alignment. Third, we argue that modeling the correlation between speech and face motions needs to consider long-term audio context dependencies [31]. Accordingly, we do not restrict the attention scope of the encoder self-attention, thus maintaining its ability to capture long-range audio context dependencies. Fourth, transformer with the sinusoidal position encoding has weak abilities to generalize to sequence lengths longer than the ones seen during training [19, 50]. Inspired by Attention with Linear Biases (ALiBi) [50], we add a temporal bias to the query-key attention score and design a periodic positional encoding strategy to improve the model’s generalization to longer audio sequences.

The main contributions of our work are as follows:

- **An autoregressive transformer-based architecture for speech-driven 3D facial animation.** FaceFormer encodes the long-term audio context and the history of face motions to autoregressively predict a sequence of

animated 3D face meshes. It achieves highly realistic and temporally stable animation of the whole face including both the upper face and the lower face.

- **The biased attention modules and a periodic positional encoding strategy.** We carefully design the biased cross-modal MH attention to align the different modalities, and the biased causal MH self-attention with a periodic positional encoding strategy to improve the generalization to longer audio sequences.
- **Effective utilization of the self-supervised pre-trained speech model.** Incorporating the self-supervised pre-trained speech model in our end-to-end architecture can not only handle the data limitation problem, but also notably improve the accuracy of mouth movements for the difficult cases, *e.g.*, the lips are fully closed on /b/,/m/,/p/ phonemes.
- **Extensive experiments and the user study to assess the quality of synthesized face motions.** The results demonstrate the superiority of FaceFormer over existing state-of-the-art methods in terms of realistic facial animation and lip sync on two 3D datasets [17, 24].

2. Related Work

2.1. Speech-Driven 3D Facial Animation

Facial animation [5, 25, 33–35, 55, 62, 72] has attracted considerable attention over the years. While aware of extensive 2D-based approaches [10, 11, 16, 18, 23, 29, 49, 59, 66, 70], we focus on animating a 3D model in this work. Typically, the procedural methods [22, 41, 54, 65] establish a set of explicit rules for animating the talking mouth. For example, the dominance functions [41] are used to characterize the speech control parameters. The dynamic viseme model proposed by Taylor *et al.* [54] exploits the one-to-many mapping of phonemes to lip motions. Xu *et al.* [65] construct a canonical set for modeling coarticulation effects. The state-of-the-art procedural approach JALI [22] utilizes two anatomical actions to animate a 3D facial rig.

One appealing strength of the above procedural methods is the explicit control of the system to ensure the accuracy of the mouth movements. However, they require a lot of manual effort in parameter tuning. Alternatively, a wide variety of data-driven approaches [6, 17, 28, 31, 40, 48, 51, 53] has been proposed to produce 3D facial animation. Cao *et al.* [6] synthesize 3D facial animation based on the proposed Anime Graph structure and a search-based technique. The sliding window approach [53] requires the transcribed phoneme sequences as input and can re-target the output to other animation rigs. An end-to-end convolutional network elaborated by Karras *et al.* [31] leverages the linear predictive coding method to encode audio and designs a la-

tent code to disambiguate the variations in facial expression. Zhou *et al.* [71] employ a three-stage network that combines phoneme groups, landmarks and audio features to predict viseme animation curves. VOCA [17] is a speaker-independent 3D facial animation method that captures a variety of speaking styles, yet the generated face motions are mostly present in the lower face. Recently, MeshTalk [51] learns a categorical latent space that successfully disentangles audio-correlated and audio-uncorrelated face motions.

The most related works are methods [17, 31, 51] where the high-resolution 3D data are used for training and the output is represented as the high-dimensional vector in 3D vertex space. The former two models [17, 31] are trained using short audio windows, thus ignoring the long-term audio context. Despite the highly realistic facial animation achieved the latter method [51], it requires large amounts of high-fidelity 3D facial data to ensure the animation quality.

2.2. Transformers in Vision and Graphics

Transformer [58] has emerged as a strong alternative to both RNN and CNN. In contrast to RNNs that process sequence tokens recursively, transformers can attend to all tokens in the input sequence parallelly, thereby modeling the long-range contextual information effectively. Vision Transformer (ViT) [21] is the first work that explores the direct application of transformers to the task of image classification. Following ViT, some follow-up works [9, 14, 56] have been introduced to boost performance for image recognition problems. Besides, transformer-based models and the variants have also been proposed in object detection [7], semantic segmentation [64], image generation [30], *etc.* In computer graphics, transformers have been exploited for 3D point cloud representations and 3D mesh, such as Point Transformer [68], Point Cloud Transformer [26] and Mesh Transformer [38]. We refer readers to the comprehensive survey [32] for further information.

Some of the most recent works on 3D body motion synthesis [1, 4, 46] and 3D dance generation [36, 37, 57] have explored the power of transformer in modeling sequential data and produced impressive results. Different from dance generation where the output motion is highly unconstrained, the task of speech-driven 3D facial animation inherently requires the alignment between audio and face motions to ensure the accuracy of lip motions. Meanwhile, the long-term audio context is expected to be considered, which is important for animating the whole face. Consequently, we present FaceFormer that incorporates the desirable properties for the speech-driven 3D facial animation problem.

3. Our Approach: FaceFormer

We formulate speech-driven 3D facial animation as a sequence-to-sequence (seq2seq) learning problem and propose a novel seq2seq architecture (Fig. 2) to autoregres-

sively predict facial movements conditioned on both audio context and past facial movement sequence. Suppose that there is a sequence of ground-truth 3D face movements $\mathbf{Y}_T = (\mathbf{y}_1, \dots, \mathbf{y}_T)$, where T is the number of visual frames, and the corresponding raw audio \mathcal{X} . The goal here is to produce a model that can synthesize facial movements $\hat{\mathbf{Y}}_T$ that is similar to \mathbf{Y}_T given the raw audio \mathcal{X} . In the encoder-decoder framework (Fig. 2), the encoder first transforms \mathcal{X} into speech representations $\mathbf{A}_{T'} = (\mathbf{a}_1, \dots, \mathbf{a}_{T'})$, where T' is the frame length of speech representations. The style embedding layer contains a set of learnable embeddings that represents speaker identities $\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_N)$. Then, the decoder autoregressively predicts facial movements $\hat{\mathbf{Y}}_T = (\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_T)$ conditioned on $\mathbf{A}_{T'}$, the style embedding \mathbf{s}_n of speaker \mathbf{n} , and the past facial movements. Formally,

$$\hat{\mathbf{y}}_t = \text{FaceFormer}_\theta(\hat{\mathbf{y}}_{<t}, \mathbf{s}_n, \mathcal{X}), \quad (1)$$

where θ denotes the model parameters, t is the current time-step in the sequence and $\hat{\mathbf{y}}_t \in \hat{\mathbf{Y}}_T$. For the remainder of this section, we describe each component of the FaceFormer architecture in detail.

3.1. FaceFormer Encoder

3.1.1 Self-Supervised Pre-Trained Speech Model

The design of our FaceFormer encoder follows the state-of-the-art self-supervised pre-trained speech model, wav2vec 2.0 [2]. Specifically, the encoder is composed of an audio feature extractor and a multi-layer transformer encoder [58]. The audio feature extractor, which consists of several temporal convolutions layers (TCN), transforms the raw waveform input into feature vectors with frequency f_a . The transformer encoder is a stack of multi-head self-attention and feed-forward layers, converting the audio feature vectors into contextualized speech representations. The outputs of the temporal convolutions are discretized to a finite set of speech units via a quantization module. Similar to masked language modeling [20], wav2vec 2.0 uses the context surrounding a masked time step to identify the true quantized speech unit by solving a contrastive task.

We initialize our encoder (Fig. 2) with the pre-trained wav2vec 2.0 weights, and add a randomly initialized linear projection layer on the top. Since the facial motion data might be captured with a frequency f_m that is different to f_a (*e.g.*, $f_a = 49\text{Hz}$ while for the BIWI dataset [24] $f_m = 25\text{fps}$), we add a linear interpolation layer after the temporal convolutions for resampling the audio features, which results in the output length kT , where $k = \lceil \frac{f_a}{f_m} \rceil$. Therefore, the outputs of the linear projection layer can be represented as $\mathbf{A}_{kT} = (\mathbf{a}_1, \dots, \mathbf{a}_{kT})$. In this way, the audio and motion modalities can be aligned by the biased cross-modal multi-head attention (Sec. 3.2.3).

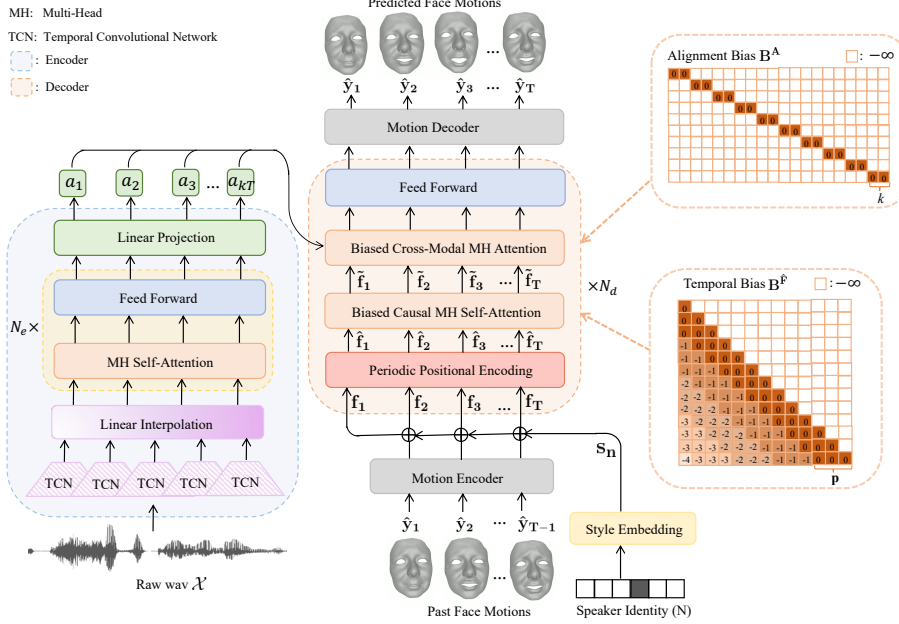


Figure 2. **Overview of FaceFormer.** An encoder-decoder model with Transformer architecture takes raw audio as input and autoregressively generates a sequence of animated 3D face meshes. Layer normalizations and residual connections are omitted for simplicity. The overall design of the FaceFormer encoder follows wav2vec 2.0 [2]. In addition, a linear interpolation layer is added after TCN for resampling the audio features. We initialize the encoder with the corresponding pre-trained wav2vec 2.0 weights. The FaceFormer decoder consists of two main modules: a biased causal MH self-attention with a periodical positional encoding for generalizing to longer input sequences, and a biased cross-modal multi-head (MH) attention for aligning audio-motion modalities. During training, the parameters of TCN are fixed, whereas the other parts of the model are learnable.

3.2. FaceFormer Decoder

3.2.1 Periodic Positional Encoding

In practice, transformer has very limited generalization abilities for longer sequences due to the sinusoidal positional encoding method [19, 50]. Attention with Linear Biases (ALiBi) [50] method is proposed to improve generalization abilities by adding a constant bias to the query-key attention score. In our experiments, we notice that directly replacing the sinusoidal positional encoding with ALiBi would lead to a static facial expression during inference. This is because ALiBi does not add any position information to the input representation, which might influence the robustness of the temporal order information, especially for our case where the training sequences have subtle motion variations among adjacent frames. To alleviate this issue, we devise a periodic positional encoding (PPE) for injecting the temporal order information, while being compatible with ALiBi. Specifically, we modify the original sinusoidal positional encoding method [58] to make it periodic with respect to a hyper-parameter p that indicates the period:

$$\begin{aligned} PPE_{(t,2i)} &= \sin\left(\frac{(t \bmod p)}{10000^{2i/d}}\right) \\ PPE_{(t,2i+1)} &= \cos\left(\frac{(t \bmod p)}{10000^{2i/d}}\right) \end{aligned} \quad (2)$$

where t denotes the token position or the current time-step in the sequence, d is the model dimension, and i is the dimension index. Rather than assigning a unique position identifier for each token [58], the proposed PPE strategy recurrently injects the position information within each period p (as shown in Section 3.2.2). Before PPE, we first project the face motion \hat{y}_t into a d -dimensional space via a motion encoder. To model the speaking style, we embed the one-hot speaker identity to a d -dimensional vector s_n via a style embedding layer and add it to the facial motion representation:

$$\mathbf{f}_t = \begin{cases} (\mathbf{W}^f \cdot \hat{y}_{t-1} + \mathbf{b}^f) + \mathbf{s}_n, & 1 < t \leq T, \\ \mathbf{s}_n, & t = 1, \end{cases} \quad (3)$$

where \mathbf{W}^f is the weight, \mathbf{b}^f is the bias and \hat{y}_{t-1} is the prediction from the last time step. Then PPE is applied to \mathbf{f}_t to provide the temporal order information periodically:

$$\hat{\mathbf{f}}_t = \mathbf{f}_t + PPE(t). \quad (4)$$

3.2.2 Biased Causal Multi-Head Self-Attention

We design a biased causal multi-head (MH) self-attention mechanism based on ALiBi [50], which is reported to be beneficial for generalizing to longer sequences in language

modeling. Given the temporally encoded facial motion representation sequence $\hat{\mathbf{F}}_t = (\hat{\mathbf{f}}_1, \dots, \hat{\mathbf{f}}_t)$, biased causal MH self-attention first linearly projects $\hat{\mathbf{F}}_t$ into queries $\mathbf{Q}^{\hat{\mathbf{F}}}$ and keys $\mathbf{K}^{\hat{\mathbf{F}}}$ of dimension d_k , and values $\mathbf{V}^{\hat{\mathbf{F}}}$ of dimension d_v . To learn the dependencies between each frame in the context of the past facial motion sequence, a weighted contextual representation is calculated by performing the scaled dot-product attention [58]:

$$\text{Att}(\mathbf{Q}^{\hat{\mathbf{F}}}, \mathbf{K}^{\hat{\mathbf{F}}}, \mathbf{V}^{\hat{\mathbf{F}}}, \mathbf{B}^{\hat{\mathbf{F}}}) = \text{softmax} \left(\frac{\mathbf{Q}^{\hat{\mathbf{F}}}(\mathbf{K}^{\hat{\mathbf{F}}})^T}{\sqrt{d_k}} + \mathbf{B}^{\hat{\mathbf{F}}} \right) \mathbf{V}^{\hat{\mathbf{F}}}, \quad (5)$$

where $\mathbf{B}^{\hat{\mathbf{F}}}$ is the temporal bias we add to ensure causality and to improve the ability to generalize to longer sequences.

More specifically, $\mathbf{B}^{\hat{\mathbf{F}}}$ is a matrix that has negative infinity in the upper triangle to avoid looking at future frames to make current predictions. For the generalization ability, we add static and non-learned biases to the lower triangle of $\mathbf{B}^{\hat{\mathbf{F}}}$. Different from ALiBi [50], we introduce the period \mathbf{p} and inject the temporal bias to each period $([1 : \mathbf{p}], [\mathbf{p} + 1 : 2\mathbf{p}], \dots)$. Let us define i and j as the indices of $\mathbf{B}^{\hat{\mathbf{F}}}$ ($1 \leq i \leq t, 1 \leq j \leq t$). Then the temporal bias $\mathbf{B}^{\hat{\mathbf{F}}}$ is formulated as:

$$\mathbf{B}^{\hat{\mathbf{F}}}(i, j) = \begin{cases} \lfloor (i - j) / \mathbf{p} \rfloor, & j \leq i, \\ -\infty, & \text{otherwise.} \end{cases} \quad (6)$$

In this way, we bias the casual attention by assigning higher attention weights to the closer period. Intuitively, the closest facial frames period $(\hat{\mathbf{y}}_{t-\mathbf{p}}, \dots, \hat{\mathbf{y}}_{t-1})$ are most likely to affect the current prediction of $\hat{\mathbf{y}}_t$. Thus, our proposed temporal bias can be considered as a generalized form of ALiBi and ALiBi becomes a special case when $\mathbf{p} = 1$.

The MH attention mechanism, which consists of \mathbf{H} parallel scaled dot-product attentions, is applied to jointly extract the complementary information from multiple representation subspaces. The outputs of \mathbf{H} heads are concatenated together and projected forward by a parameter matrix $\mathbf{W}^{\hat{\mathbf{F}}}$:

$$\text{MH}(\mathbf{Q}^{\hat{\mathbf{F}}}, \mathbf{K}^{\hat{\mathbf{F}}}, \mathbf{V}^{\hat{\mathbf{F}}}, \mathbf{B}^{\hat{\mathbf{F}}}) = \text{Concat}(\text{head}_1, \dots, \text{head}_{\mathbf{H}}) \mathbf{W}^{\hat{\mathbf{F}}},$$

where $\text{head}_{\mathbf{h}} = \text{Att}(\mathbf{Q}_{\mathbf{h}}^{\hat{\mathbf{F}}}, \mathbf{K}_{\mathbf{h}}^{\hat{\mathbf{F}}}, \mathbf{V}_{\mathbf{h}}^{\hat{\mathbf{F}}}, \mathbf{B}_{\mathbf{h}}^{\hat{\mathbf{F}}})$. (7)

Similar to ALiBi [50], we add a head-specific scalar \mathbf{m} for the MH setting. For each head $_{\mathbf{h}}$, the temporal bias is defined as $\mathbf{B}_{\mathbf{h}}^{\hat{\mathbf{F}}} = \mathbf{B}^{\hat{\mathbf{F}}} \cdot \mathbf{m}$. The scalar \mathbf{m} is a head-specific slope and is not learned during training. For \mathbf{H} heads, \mathbf{m} will start at $2^{-2^{(-\log_2 \mathbf{H} + 3)}}$ and multiply each element by the same value to compute the next element. Concretely, if the model has 4 heads, the corresponding slopes will be 2^{-2} , 2^{-4} , 2^{-6} and 2^{-8} .

3.2.3 Biased Cross-Modal Multi-Head Attention

The biased cross-modal multi-head attention aims to combine the outputs of Faceformer encoder (speech features) and biased causal MH self-attention (motion features) to align the audio and motion modalities (see Fig. 2). For this purpose, we add an alignment bias to the query-key attention score, which is simple and effective. The alignment bias $\mathbf{B}^{\mathbf{A}}$ ($1 \leq i \leq t, 1 \leq j \leq k\mathbf{T}$) is represented as:

$$\mathbf{B}^{\mathbf{A}}(i, j) = \begin{cases} 0, & ki \leq j < k(i + 1) \\ -\infty, & \text{otherwise} \end{cases}, \quad (8)$$

Each token in $\mathbf{A}_{k\mathbf{T}}$ has captured the long-term audio context due to the self-attention mechanism. On the other hand, assuming the outputs of biased causal MH self-attention is $\tilde{\mathbf{F}}_t = (\tilde{\mathbf{f}}_1, \dots, \tilde{\mathbf{f}}_t)$, each token in $\tilde{\mathbf{F}}_t$ has encoded the history context of face motions. Both $\mathbf{A}_{k\mathbf{T}}$ and $\tilde{\mathbf{F}}_t$ are fed into biased cross-modal MH attention. Likewise, $\mathbf{A}_{k\mathbf{T}}$ is transformed into two separate matrices: keys $\mathbf{K}^{\mathbf{A}}$ and values $\mathbf{V}^{\mathbf{A}}$, whereas $\tilde{\mathbf{F}}_t$ is transformed into queries $\mathbf{Q}^{\tilde{\mathbf{F}}}$. The output is calculated as a weighted sum of $\mathbf{V}^{\mathbf{A}}$,

$$\text{Att}(\mathbf{Q}^{\tilde{\mathbf{F}}}, \mathbf{K}^{\mathbf{A}}, \mathbf{V}^{\mathbf{A}}, \mathbf{B}^{\mathbf{A}}) = \text{softmax} \left(\frac{\mathbf{Q}^{\tilde{\mathbf{F}}}(\mathbf{K}^{\mathbf{A}})^T}{\sqrt{d_k}} + \mathbf{B}^{\mathbf{A}} \right) \mathbf{V}^{\mathbf{A}}. \quad (9)$$

To explore different subspaces, we also extend Eq. (9) to \mathbf{H} heads as in Eq. (7). Finally, the predicted face motion $\hat{\mathbf{y}}_t$ is obtained by projecting the d -dimensional hidden state back to the \mathbf{V} -dimensional 3D vertex space via a motion decoder.

3.3. Training and Testing

During the training phase, we adopt an autoregressive scheme instead of a teacher-forcing scheme. In our experiments, we observe that training FaceFormer with a less guided scheme works better than a fully guided one. Once the complete 3D facial motion sequence is produced, the model is trained by minimizing the Mean Squared Error (MSE) between the decoder outputs $\hat{\mathbf{Y}}_t = (\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_{\mathbf{T}})$ and the ground truth $\mathbf{Y}_t = (\mathbf{y}_1, \dots, \mathbf{y}_{\mathbf{T}})$:

$$\mathcal{L}_{\text{MSE}} = \sum_{t=1}^{\mathbf{T}} \sum_{v=1}^{\mathbf{V}} \|\hat{\mathbf{y}}_{t,v} - \mathbf{y}_{t,v}\|^2, \quad (10)$$

where \mathbf{V} represents the number of vertices of the 3D face mesh.

At inference time, FaceFormer autoregressively predicts a sequence of animated 3D face meshes. More specifically, at each time-step, it predicts the face motion $\hat{\mathbf{y}}_t$ conditioned on the raw audio \mathcal{X} , the history of face motions $\hat{\mathbf{y}}_{<t}$ and the style representations \mathbf{s}_n as in Eq. (1). \mathbf{s}_n is determined by the speaker identity, and thus altering the one-hot identity vector can manipulate the output in different styles.

4. Experiments and Results

4.1. Experimental Settings

We use two publicly available 3D datasets, BIWI [24] and VOCASET [17] for training and testing. Both datasets provide the audio-3D scan pairs of English spoken utterances. BIWI contains 40 unique sentences shared across all speakers. VOCASET contains 255 unique sentences, some of which are shared across speakers. Comparatively, BIWI represents a more challenging dataset for lip sync as it covers fewer phonemes.

BIWI Dataset. BIWI is a corpus of affective speech and corresponding dense dynamic 3D face geometries. 14 human subjects are asked to read 40 English sentences, each of which is recorded twice: in a neutral or emotional context. The 3D face geometries are captured at 25fps, each with 23370 vertices. Each sequence is 4.67 seconds long on average. For our experiments, we use the subset where the sentences are recorded in the emotional context. Specifically, we split the data into a training set (BIWI-Train) of 192 sentences spoken by six subjects (each subject speaks 32 sentences), a validation set (BIWI-Val) of 24 sentences spoken by six subjects (each subject speaks 4 sentences), and two testing sets (BIWI-Test-A and BIWI-Test-B). BIWI-Test-A includes 24 sentences spoken by six seen subjects (each speaks 4 sentences), and BIWI-Test-B includes 32 sentences spoken by eight unseen subjects (each speaks 4 sentences).

VOCASET Dataset. VOCASET is composed of 480 facial motion sequences from 12 subjects. Each sequence is captured at 60fps and is between 3 and 4 seconds long. Each 3D face mesh has 5023 vertices. For a fair comparison, we use the same training, validation and testing splits as VOCA [17], which we refer to VOCA-Train, VOCA-Val and VOCA-Test, respectively.

Baseline Methods. We compare FaceFormer with two state-of-the-art methods, VOCA [17] and MeshTalk [51], on both BIWI and VOCASET. Among the three methods, FaceFormer and VOCA require conditioning on a training speaker identity during inference. For unseen subjects, we obtain the predictions of FaceFormer and VOCA by conditioning on all training identities. The implementation details of FaceFormer and the baseline methods are provided in the supplementary material (Sec. 1 and Sec. 2).

4.2. Evaluation Results

Lip-sync Evaluation. We follow the lip-sync metric employed in MeshTalk [51] for evaluating the quality of lip movements. The maximal L2 error of all lip vertices is defined as the lip error for each frame. The error is calculated by comparing the predictions and the processed 3D face geometry data. We report the computed average over all testing sequences of BIWI-test-A for VOCA, MeshTalk and

Table 1. Comparison of lip-sync errors. We compare FaceFormer with two state-of-the-art methods [17, 51] on BIWI-Test-A. The maximal lip vertex error [51] is used for lip-sync evaluation.

Methods	Lip Vertex Error ($\times 10^{-4}$ mm)
VOCA	7.6427
MeshTalk	6.7436
FaceFormer (Ours)	5.3742

Table 2. User study results on BIWI-Test-B. We use A/B testing and report the percentage of answers where A is preferred over B.

Ours vs. Competitor	Realism	Lip Sync
Ours vs. VOCA	83.85 \pm 3.76	82.64 \pm 3.77
Ours vs. MeshTalk	83.33 \pm 4.07	80.56 \pm 5.22
Ours vs. GT	35.24 \pm 2.87	36.98 \pm 1.38

Table 3. User study results on VOCA-Test.

Ours vs. Competitor	Realism	Lip Sync
Ours vs. VOCA	77.92 \pm 7.94	77.08 \pm 7.32
Ours vs. MeshTalk	82.92 \pm 2.60	82.08 \pm 3.15
Ours vs. GT	29.17 \pm 10.41	30.42 \pm 8.04

FaceFormer in Tab. 1. The lower average lip error achieved by FaceFormer suggests it can produce more accurate lip movements compared to the other two methods.

Qualitative Evaluation. Given the many-to-many mappings between upper face motions and the speech utterance, it is suggested that qualitative evaluations and user studies are more proper for evaluating the quality of speech-driven facial animation than using quantitative metrics [17, 31]. We refer the readers to our supplementary video for the assessment of the motion quality. The video compares the results of our approach, those by the previous methods [17, 31, 51, 53] and the ground truth. Specifically, we test our model using (1) audio sequences from BIWI and VOCASET test sets, (2) audio clips extracted from supplementary videos of previous methods and (3) audio clips extracted from TED videos on YouTube. For the last two cases, the results are predicted from the model trained on BIWI. The video shows that FaceFormer produces realistic and natural-looking facial animation with accurate lip synchronization. Compared to VOCA and MeshTalk, it is notable that, FaceFormer produces more realistic facial motions and better lip sync with proper mouth closures in many situations, *e.g.*, the lips are fully closed when pronouncing /b/, /m/, /p/. We also show that our system can produce animation of talking in different styles and different languages.

4.3. Perceptual Evaluation

User Study on BIWI. We conduct user studies on Amazon Mechanical Turk (AMT) to evaluate the animation quality of FaceFormer, compared with the ground truth, VOCA and MeshTalk. For BIWI, we obtain the results of three methods using all test audio sequences of BIWI-Test-B (32 sentences). The results of FaceFormer and VOCA are

produced by conditioning on all training speaker identities, which results in 192 videos (32 sentences \times 6 identities) for each method. Therefore, 576 A vs. B pairs (192 videos \times 3 comparisons) are created for BIWI-Test-B. For each HIT (human intelligence task), the AMT interface shows four video pairs including the qualification test in randomized order, and the Turker is instructed to judge the videos in terms of realistic facial animation and lip sync. Each video pair is evaluated by three Turkers. In particular, Turkers must pass the qualification test otherwise they are not allowed to submit HITs. Finally, we collect 576 HITs for the user study on BIWI. More details about the user study are described in our supplementary material (Sec. 3). Tab. 2 shows the percentage of A/B testing in terms of realism and lip sync. Turkers favor FaceFormer over VOCA in terms of realistic facial animation and lip sync. We believe this is mainly due to two reasons: (1) face motions synthesized by VOCA are mostly present in the lower face; (2) VOCA sometimes fails to fully close the mouth at the phonemes /b/, /m/, /p/. FaceFormer also outperforms MeshTalk and we attribute this to the results produced by FaceFormer having more expressive facial motions and more accurate mouth movements. Not surprisingly, Turkers perceive the ground truth more realistic than FaceFormer.

User Study on VOCASET. In the second user study, we compare the results of three methods on VOCA-Test. We randomly select 10 sentences from VOCA-Test and obtain the results of FaceFormer and VOCA conditioned on all training speaker identities, which results in 80 videos (10 sentences \times 8 identities) for each method. In total, 240 A vs. B pairs (80 videos \times 3 comparisons) are created for VOCA-Test. Similarly, for each pair, Turkers make the choice between two videos in terms of realism and lip sync. Since VOCASET has very few upper face motions, movements are present mostly in the lower face for all three methods. In this case, well-synchronized lip motions are important for generating perceptually realistic results. Tab. 3 shows that FaceFormer achieves higher percentages over VOCA and MeshTalk. We believe this is because our results have better synchronized mouth shapes and closures. Similarly, there is still a certain gap between our results and the ground truth.

4.4. Visualization Analysis

To provide insights into the underlying attention mechanism, we visualize the attention weights for the MH self-attention of the encoder, as well as the biased causal MH self-attention of the decoder. We consider 100 frames of a test sequence from BIWI and examine the attention weights that are used to predict the last frame. Fig. 3 visualizes the average attention weights across all heads. We observe that the encoder self-attention (Fig. 3 (a)) not only focuses on the nearby audio frames (as reflected by the diagonal line)




Figure 3. **Attention Weights Visualization.** Attention weights of the (a) MH self-attention of the encoder and (b) biased causal MH self-attention of the decoder.

but also attends to some farther future and past frames. This indicates that the self-attention mechanism of transformer is able to capture both the short- and long-range audio context dependencies. The attended audio frames may contain more informative context features that influence the current face motion. For the decoder self-attention (Fig. 3 (b)), the visualization corresponds to the casual attention incorporated with the temporal bias (Eq. (6)). There is a clear pattern that the face motion frames in a closer period are assigned with higher weights, as those frames are more likely to influence the current face motion. For example, there is a high probability that people will keep smiling if they have been smiling over the past frames.

4.5. Ablation Study

The visual results of ablation are included in the supplementary video. Please watch the video for comparison.

4.5.1 Ablation on FaceFormer Encoder

Effect of the encoder self-attention module. To investigate the effect of the MH self-attention module in FaceFormer encoder, we directly remove it from the whole architecture, with the pre-trained TCN retained to extract the speech representations. We refer to this variant as “TCN+FaceFormer Decoder” and conduct the comparison experiments on BIWI. The results show that “TCN+FaceFormer Decoder” often fails to close the mouth, resulting in out-of-sync lip motions. Besides, the produced results have a temporal jitter effect around the mouth region, as shown in the supplementary video.

Effect of the wav2vec weights initialization. We also perform an ablation study of the wav2vec weights initialization by comparing FaceFormer trained with and without wav2vec weights initialization (denoted as “FaceFormer w/o wav2vec”). Without wav2vec weights initialization, we observe a downgrade of the quality of face movements. “FaceFormer w/o wav2vec” can not produce synchronized mouth motions and a temporal jitter effect can be observed. This suggests that simply training FaceFormer with randomly initialized weights might converge to a poor solution.

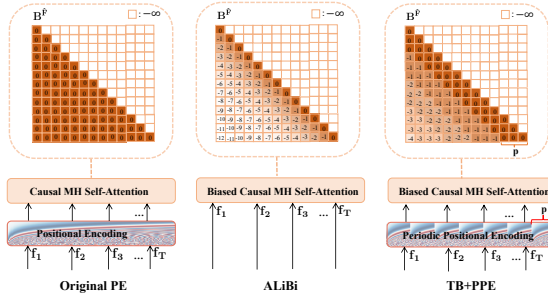


Figure 4. Illustration of different positional encoding strategies.

Hence, the wav2vec weights initialization is necessary for the FaceFormer encoder.

4.5.2 Ablation on FaceFormer Decoder

Choices of the decoder architecture. We explore whether the transformer-based architecture has advantages over a FC layer or LSTM by training and testing two alternative variants: “FaceFormer Encoder+FC” and “FaceFormer Encoder+LSTM”. As shown in the video, FaceFormer yields more stable mouth motions and more accurate lip sync. Compared to the FC decoder, the autoregressive mechanism of FaceFormer decoder can stabilize the lip motions by modeling the history motions. On the other hand, the self-attention mechanism of FaceFormer decoder might model the context cues in history motions better than LSTM, thus having more temporally coherent lip motions.

Effect of the alignment bias. We examine the effect of the alignment bias (Eq. (8)) by removing it from the biased cross-modal MH attention module. The model without the alignment bias (denoted as “FaceFormer w/o AB”) tends to generate muted facial expressions across all frames. Hence, the alignment bias is indispensable for the cross-modal attention in aligning the audio-motion modalities correctly.

Effect of the proposed positional encoding strategy. The proposed positional encoding strategy is adding a temporal bias to the attention score and making the original sinusoidal position embedding [58] (“Original PE”) periodic. We refer to this strategy as “TB+PPE” and compare it with “Original PE” and “ALiBi” [50]. The differences of three positional encoding strategies are visualized in Fig. 4. A detailed illustration of “PPE” can be found in our supplementary material (Sec. 4). The results show “Original PE” produces well-synchronized mouth motions with proper lip closures, yet has a temporal jitter effect around the lips during silent frames, especially as the test audio sequence exceeds the average length of training audio sequences. While “ALiBi” does not influence the results on BIWI, it quickly freezes to a static facial expression when training and testing on VOCASET. This happens because the ALiBi does not add any position information to the input representation,

which might influence the robustness of the temporal order information. This influence is more obvious when motion data have subtle variations among adjacent frames. Since “TB+PPE” is a key component for improving the ability to generalize to longer sequences, we additionally study its influence by conducting the perceptual evaluation on AMT. Specifically, we download the TED videos shared under the “CC BY-NC-ND 4.0 International License”, and extract 15 representative audio clips for the user study. The audio sequences are around 20 seconds long, more than four times the average length of training audio sequences. For the comparison to “Original PE”, we randomly sample a training identity and use it as the condition for both methods. Similar to the user study in Sec. 4.3, each video pair is evaluated by three judges. Overall, Turkers perceive the facial animation results of FaceFormer more realistic ($57.78\% \pm 16.78\%$) and the generated lip motions of FaceFormer more in sync with audio ($62.22\% \pm 10.18\%$) than “Original PE”. The result indicates FaceFormer generalizes better to longer audio clips. The likely explanation is that “Original PE” tends to generate unstable lip motions during silent frames when testing on longer audio sequences.

5. Discussion and Conclusion

In this work, we propose an autoregressive transformer-based architecture for speech-driven 3D facial animation. The encoder effectively leverages the self-supervised pre-trained speech representations, and the inside self-attention captures long-range audio context dependencies. The decoder attention modules with periodic positional encoding are tailored for cross-modal alignment and generalization to longer sequences. Overall, FaceFormer demonstrates higher quality for lip synchronization and realistic facial animation compared to the state-of-the-arts. However, our model requires access to the full audio sequence, making it not suitable for online streaming applications. Another bottleneck is the quadratic memory and time complexity of self-attention. One future work is to address this problem using advanced techniques [3, 61] that improve the efficiency of self-attention.

Ethics Considerations: We should use technology responsibly and be careful about the synthesized content. Since our technique requires 3D scan data collected from actors, it is important to obtain consent from the actors during data acquisition. Our method can animate a realistic 3D talking face from an arbitrary audio signal. However, there is a risk that such techniques could potentially be misused to cause embarrassment. Thus, we hope to raise the public’s awareness about the risks of the potential misuse and encourage research efforts on the responsible use of technology.

Acknowledgement. This research is partly supported by New Energy and Industrial Technology Development Organization (NEDO) (ref:JPNP21004).

References

- [1] Emre Aksan, Peng Cao, Manuel Kaufmann, and Otmar Hilliges. A spatio-temporal transformer for 3d human motion prediction. *arXiv preprint arXiv:2004.08692*, 2020. [2](#), [3](#)
- [2] Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. *arXiv preprint arXiv:2006.11477*, 2020. [2](#), [3](#), [4](#)
- [3] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020. [8](#)
- [4] Uttaran Bhattacharya, Nicholas Rewkowski, Abhishek Banerjee, Pooja Guhan, Aniket Bera, and Dinesh Manocha. Text2gestures: A transformer-based network for generating emotive body gestures for virtual agents. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*, pages 1–10. IEEE, 2021. [2](#), [3](#)
- [5] Chen Cao, Hongzhi Wu, Yanlin Weng, Tianjia Shao, and Kun Zhou. Real-time facial animation with image-based dynamic avatars. *ACM Transactions on Graphics*, 35(4), 2016. [2](#)
- [6] Yong Cao, Wen C Tien, Petros Faloutsos, and Frédéric Pighin. Expressive speech-driven facial animation. *ACM Transactions on Graphics*, 24(4):1283–1302, 2005. [2](#)
- [7] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Proceedings of the European Conference on Computer Vision*, pages 213–229. Springer, 2020. [3](#)
- [8] Yujin Chai, Yanlin Weng, Lvdi Wang, and Kun Zhou. Speech-driven facial animation with spectral gathering and temporal attention. *Frontiers of Computer Science*, 2020. [1](#)
- [9] Chun-Fu Chen, Quanfu Fan, and Rameswar Panda. Crossvit: Cross-attention multi-scale vision transformer for image classification. *arXiv preprint arXiv:2103.14899*, 2021. [3](#)
- [10] Lele Chen, Guofeng Cui, Celong Liu, Zhong Li, Ziyi Kou, Yi Xu, and Chenliang Xu. Talking-head generation with rhythmic head motion. In *Proceedings of the European Conference on Computer Vision*, pages 35–51. Springer, 2020. [2](#)
- [11] Lele Chen, Zhiheng Li, Ross K Maddox, Zhiyao Duan, and Chenliang Xu. Lip movements generation at a glance. In *Proceedings of the European Conference on Computer Vision*, pages 520–535, 2018. [2](#)
- [12] 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 Conference on Computer Vision and Pattern Recognition*, pages 7832–7841, 2019. [1](#)
- [13] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pre-training from pixels. In *International Conference on Machine Learning*, pages 1691–1703, 2020. [2](#)
- [14] Xiangxiang Chu, Zhi Tian, Bo Zhang, Xinlong Wang, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Conditional positional encodings for vision transformers. *arXiv preprint arXiv:2102.10882*, 2021. [3](#)
- [15] Joon Son Chung, Amir Jamaludin, and Andrew Zisserman. You said that? *arXiv preprint arXiv:1705.02966*, 2017. [1](#)
- [16] 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. [2](#)
- [17] Daniel Cudeiro, Timo Bolkart, Cassidy Laidlaw, Anurag Ranjan, and Michael J Black. Capture, learning, and synthesis of 3d speaking styles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10101–10111, 2019. [1](#), [2](#), [3](#), [6](#)
- [18] Dipanjan Das, Sandika Biswas, Sanjana Sinha, and Brojeshwar Bhowmick. Speech-driven facial animation using cascaded gans for learning of motion and texture. In *European Conference on Computer Vision*, pages 408–424. Springer, 2020. [2](#)
- [19] Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. Universal transformers. *arXiv preprint arXiv:1807.03819*, 2018. [2](#), [4](#)
- [20] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. [2](#), [3](#)
- [21] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [2](#), [3](#)
- [22] Pif Edwards, Chris Landreth, Eugene Fiume, and Karan Singh. Jali: an animator-centric viseme model for expressive lip synchronization. *ACM Transactions on Graphics*, 35(4):1–11, 2016. [2](#)
- [23] Bo Fan, Lijuan Wang, Frank K Soong, and Lei Xie. Photoreal talking head with deep bidirectional lstm. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 4884–4888. IEEE, 2015. [2](#)
- [24] Gabriele Fanelli, Juergen Gall, Harald Romsdorfer, Thibaut Weise, and Luc Van Gool. A 3-d audio-visual corpus of affective communication. *IEEE Transactions on Multimedia*, 12(6):591–598, 2010. [2](#), [3](#), [6](#)
- [25] Ohad Fried, Ayush Tewari, Michael Zollhöfer, Adam Finkelstein, Eli Shechtman, Dan B Goldman, Kyle Genova, Zeyu Jin, Christian Theobalt, and Maneesh Agrawala. Text-based editing of talking-head video. *ACM Transactions on Graphics*, 38(4):1–14, 2019. [2](#)
- [26] Meng-Hao Guo, Jun-Xiong Cai, Zheng-Ning Liu, Tai-Jiang Mu, Ralph R Martin, and Shi-Min Hu. Pct: Point cloud transformer. *Computational Visual Media*, 7(2):187–199, 2021. [3](#)
- [27] Ikhsanul Habibie, Weipeng Xu, Dushyant Mehta, Lingjie Liu, Hans-Peter Seidel, Gerard Pons-Moll, Mohamed Elgharib, and Christian Theobalt. Learning speech-driven 3d conversational gestures from video. *arXiv preprint arXiv:2102.06837*, 2021. [1](#)
- [28] Ahmed Hussen Abdelaziz, Barry-John Theobald, Paul Dixon, Reinhard Knothe, Nicholas Apostoloff, and Sachin Kajareker. Modality dropout for improved performance-

- driven talking faces. In *Proceedings of the 2020 International Conference on Multimodal Interaction*, pages 378–386, 2020. [2](#)
- [29] 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 Conference on Computer Vision and Pattern Recognition*, pages 14080–14089, 2021. [2](#)
- [30] Yifan Jiang, Shiyu Chang, and Zhangyang Wang. Transgan: Two transformers can make one strong gan. *arXiv preprint arXiv:2102.07074*, 2021. [3](#)
- [31] Tero Karras, Timo Aila, Samuli Laine, Antti Herva, and Jaakko Lehtinen. Audio-driven facial animation by joint end-to-end learning of pose and emotion. *ACM Transactions on Graphics*, 36(4):1–12, 2017. [1](#), [2](#), [3](#), [6](#)
- [32] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *arXiv preprint arXiv:2101.01169*, 2021. [2](#), [3](#)
- [33] Hyeonwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Niessner, Patrick Pérez, Christian Richardt, Michael Zollhöfer, and Christian Theobalt. Deep video portraits. *ACM Transactions on Graphics*, 37(4):1–14, 2018. [2](#)
- [34] Avisek Lahiri, Vivek Kwatra, Christian Frueh, John Lewis, and Chris Bregler. Lipsync3d: Data-efficient learning of personalized 3d talking faces from video using pose and lighting normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2755–2764, 2021. [2](#)
- [35] Hao Li, Jihun Yu, Yuting Ye, and Chris Bregler. Realtime facial animation with on-the-fly correctives. *ACM Transactions on Graphics*, 32(4):42–1, 2013. [2](#)
- [36] Jiaman Li, Yihang Yin, Hang Chu, Yi Zhou, Tingwu Wang, Sanja Fidler, and Hao Li. Learning to generate diverse dance motions with transformer. *arXiv preprint arXiv:2008.08171*, 2020. [2](#), [3](#)
- [37] Ruilong Li, Shan Yang, David Ross, and Angjoo Kanazawa. Learn to dance with aist++: Music conditioned 3d dance generation. *arXiv preprint arXiv:2101.08779*, 2021. [2](#), [3](#)
- [38] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-end human pose and mesh reconstruction with transformers. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1954–1963, 2021. [3](#)
- [39] Jingying Liu, Binyuan Hui, Kun Li, Yunke Liu, Yu-Kun Lai, Yuxiang Zhang, Yebin Liu, and Jingyu Yang. Geometry-guided dense perspective network for speech-driven facial animation. *IEEE Transactions on Visualization and Computer Graphics*, 2021. [1](#)
- [40] Yilong Liu, Feng Xu, Jinxiang Chai, Xin Tong, Lijuan Wang, and Qiang Huo. Video-audio driven real-time facial animation. *ACM Transactions on Graphics*, 34(6):1–10, 2015. [2](#)
- [41] DW Massaro, MM Cohen, M Tabain, J Beskow, and R Clark. Animated speech: research progress and applications. *Audiovisual Speech Processing*, page 309–345, 2012. [2](#)
- [42] Gaurav Mittal and Baoyuan Wang. Animating face using disentangled audio representations. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, pages 3290–3298, 2020. [1](#)
- [43] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: an asr corpus based on public domain audio books. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 5206–5210, 2015. [2](#)
- [44] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In *International Conference on Machine Learning*, pages 4055–4064. PMLR, 2018. [2](#)
- [45] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *International Conference on Machine Learning*, pages 1310–1318, 2013. [2](#)
- [46] Mathis Petrovich, Michael J Black, and Gül Varol. Action-conditioned 3d human motion synthesis with transformer vae. *arXiv preprint arXiv:2104.05670*, 2021. [2](#), [3](#)
- [47] Hai X Pham, Samuel Cheung, and Vladimir Pavlovic. Speech-driven 3d facial animation with implicit emotional awareness: a deep learning approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 80–88, 2017. [1](#)
- [48] Hai Xuan Pham, Yuting Wang, and Vladimir Pavlovic. End-to-end learning for 3d facial animation from speech. In *Proceedings of the ACM International Conference on Multimodal Interaction*, pages 361–365, 2018. [2](#)
- [49] 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. [2](#)
- [50] Ofir Press, Noah A Smith, and Mike Lewis. Train short, test long: Attention with linear biases enables input length extrapolation. *arXiv preprint arXiv:2108.12409*, 2021. [2](#), [4](#), [5](#), [8](#)
- [51] Alexander Richard, Michael Zollhöfer, Yandong Wen, Fernando de la Torre, and Yaser Sheikh. Meshtalk: 3d face animation from speech using cross-modality disentanglement. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1173–1182, 2021. [1](#), [2](#), [3](#), [6](#)
- [52] Supasorn Suwajanakorn, Steven M Seitz, and Ira Kemelmacher-Shlizerman. Synthesizing obama: learning lip sync from audio. *ACM Transactions on Graphics*, 36(4):1–13, 2017. [1](#)
- [53] Sarah Taylor, Taehwan Kim, Yisong Yue, Moshe Mahler, James Krahe, Anastasio Garcia Rodriguez, Jessica Hodgins, and Iain Matthews. A deep learning approach for generalized speech animation. *ACM Transactions on Graphics*, 36(4):1–11, 2017. [2](#), [6](#)
- [54] Sarah L Taylor, Moshe Mahler, Barry-John Theobald, and Iain Matthews. Dynamic units of visual speech. In *Proceedings of the ACM SIGGRAPH/Eurographics conference on Computer Animation*, pages 275–284, 2012. [2](#)
- [55] Justus Thies, Mohamed Elgharib, Ayush Tewari, Christian Theobalt, and Matthias Nießner. Neural voice puppetry: Audio-driven facial reenactment. In *Proceedings of Eu-*

- ropean Conference on Computer Vision, pages 716–731. Springer, 2020. 2
- [56] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357, 2021. 3
- [57] Guillermo Valle-Pérez, Gustav Eje Henter, Jonas Beskow, André Holzapfel, Pierre-Yves Oudeyer, and Simon Alexanderson. Transflower: probabilistic autoregressive dance generation with multimodal attention. *arXiv preprint arXiv:2106.13871*, 2021. 2, 3
- [58] 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, 3, 4, 5, 8
- [59] Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Realistic speech-driven facial animation with gans. *International Journal of Computer Vision*, 128(5):1398–1413, 2020. 2
- [60] Qianyun Wang, Zhenfeng Fan, and Shihong Xia. 3d-talkemo: Learning to synthesize 3d emotional talking head. *arXiv preprint arXiv:2104.12051*, 2021. 1
- [61] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020. 8
- [62] Thibaut Weise, Sofien Bouaziz, Hao Li, and Mark Pauly. Realtime performance-based facial animation. *ACM Transactions on Graphics*, 30(4):1–10, 2011. 2
- [63] Olivia Wiles, A Koepke, and Andrew Zisserman. X2face: A network for controlling face generation using images, audio, and pose codes. In *Proceedings of the European Conference on Computer Vision*, pages 670–686, 2018. 1
- [64] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *arXiv preprint arXiv:2105.15203*, 2021. 3
- [65] Yuyu Xu, Andrew W Feng, Stacy Marsella, and Ari Shapiro. A practical and configurable lip sync method for games. In *Proceedings of Motion on Games*, pages 131–140, 2013. 2
- [66] Ran Yi, Zipeng Ye, Juyong Zhang, Hujun Bao, and Yong-Jin Liu. Audio-driven talking face video generation with learning-based personalized head pose. *arXiv preprint arXiv:2002.10137*, 2020. 2
- [67] Dan Zeng, Han Liu, Hui Lin, and Shiming Ge. Talking face generation with expression-tailored generative adversarial network. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 1716–1724, 2020. 1
- [68] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip Torr, and Vladlen Koltun. Point transformer. *arXiv preprint arXiv:2012.09164*, 2020. 3
- [69] Hang Zhou, Yu Liu, Ziwei Liu, Ping Luo, and Xiaogang Wang. Talking face generation by adversarially disentangled audio-visual representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9299–9306, 2019. 1
- [70] 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 Conference on Computer Vision and Pattern Recognition*, pages 4176–4186, 2021. 2
- [71] Yang Zhou, Zhan Xu, Chris Landreth, Evangelos Kalogerakis, Subhansu Maji, and Karan Singh. Visemenet: Audio-driven animator-centric speech animation. *ACM Transactions on Graphics*, 37(4):1–10, 2018. 3
- [72] Michael Zollhöfer, Justus Thies, Pablo Garrido, Derek Bradley, Thabo Beeler, Patrick Pérez, Marc Stamminger, Matthias Nießner, and Christian Theobalt. State of the art on monocular 3d face reconstruction, tracking, and applications. In *Computer Graphics Forum*, pages 523–550, 2018. 2