

# StreamingTalker: Audio-driven 3D Facial Animation with Autoregressive Diffusion Model

Yifan Yang<sup>1,2</sup>, Zhi Cen<sup>1</sup>, Sida Peng<sup>1</sup>, Xiangwei Chen<sup>3</sup>,  
Yifu Deng<sup>2</sup>, Xinyu Zhu<sup>2</sup>, Fan Jia<sup>2</sup>, Xiaowei Zhou<sup>1</sup>, Hujun Bao<sup>1\*</sup>

<sup>1</sup>State Key Laboratory of CAD&CG, Zhejiang University

<sup>2</sup>Ant Group

<sup>3</sup>College of Computer Science, Zhejiang University

yfyang@zju.edu.cn

## Abstract

This paper focuses on the task of speech-driven 3D facial animation, which aims to generate realistic and synchronized facial motions driven by speech inputs. Recent methods have employed audio-conditioned diffusion models for 3D facial animation, achieving impressive results in generating expressive and natural animations. However, these methods process the whole audio sequences in a single pass, which poses two major challenges: they tend to perform poorly when handling audio sequences that exceed the training horizon and will suffer from significant latency when processing long audio inputs. To address these limitations, we propose a novel autoregressive diffusion model that outputs facial motions in a streaming manner. This design ensures flexibility with varying audio lengths and achieves low latency independent of audio duration. Specifically, we select a limited number of past frames as historical motion context and combine them with the audio input to create a dynamic condition. This condition guides a lightweight diffusion head to iteratively generate facial motion frames, enabling real-time synthesis with high-quality results. Experiments conducted on public datasets demonstrate that our approach outperforms recent baseline methods.

**Code** — <https://zju3dv.github.io/StreamingTalker/>

## Introduction

Speech-driven 3D facial animation has emerged as a critical area of research, garnering significant attention due to its wide-ranging applications in virtual reality, gaming, and telecommunication. This technology enables lifelike digital humans that can interact naturally with users through speech, enhancing both immersion and accessibility. The goal of this work is to generate a 3D facial mesh that accurately reflects the given speech audio, ensuring precise lip-synchronization and expressive motion. However, this task remains challenging, as it demands not only high visual fidelity in facial expressions but also efficient processing to support real-time generation.

Traditional approaches (Massaro et al. 12; Edwards et al. 2016) to speech-driven 3D facial animation have predominantly employed deterministic methods that directly map

phonemes to the corresponding 3D facial mesh. However, these methods often rely heavily on intermediate representations of phonemes and primarily focus on modeling mouth movements, neglecting other facial dynamics. Moreover, these traditional methods lack the capability for one-to-many generations, which significantly limits their ability to capture the variability and expressiveness of natural facial motion.

In recent years, diffusion models have achieved remarkable success in image generation tasks (Ramesh et al. 2021), leading to their introduction in the domain of speech-driven facial animation (Ma et al. 2024; Sun et al. 2024). Notable methods such as FaceDiffuser (Stan, Haque, and Yumak 2023) have demonstrated the potential of diffusion processes in this context, achieving fast denoising and good results on short sequences. However, we observe that these designs exhibit a performance drop when generating motions that extend beyond the training horizon. This limitation arises because they are trained on fixed-length sequences, which hinders their ability to generalize to longer and more complex sequences. Furthermore, as these methods employ large denoising networks such as Transformers or U-Nets and need to process the entire sequence before producing results, they incur heavy computation and struggle to support real-time applications.

Several critical obstacles must be addressed to enable real-time generation of arbitrary-length facial meshes. First, it is essential to develop a mechanism that dynamically integrates historical information from past motions, allowing the model to handle streaming outputs effectively. Additionally, maintaining a balance between generation quality and inference speed remains a key difficulty.

To tackle these challenges, we propose a novel approach that leverages an autoregressive (AR) diffusion model. Specifically, our method reformulates full-sequence generation as an AR diffusion process conditioned on historical motion. By encoding a limited number of past frames and the current audio input with an AR transformer, we derive a dynamic condition that guides a lightweight MLP-based diffusion head to generate facial meshes in real time. This approach offers several key advantages, including the ability to enable the flexible and dynamic use of past motions which enhances the model’s generalization capabilities, and the facilitation of real time generation of facial meshes for

\*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

arbitrary-length sequences, supporting streaming output and ensuring seamless rendering.

Experimental results demonstrate that our approach outperforms state-of-the-art methods on two benchmark datasets, establishing its effectiveness and robustness. Furthermore, our method achieves superior performance in long-sequence generation and greatly improves inference latency compared with previous diffusion-based models.

## Related Work

### Speech-driven 3D Facial Animation

Existing approaches to speech-driven 3D facial animation can be broadly categorized into linguistics-based and learning-based methods. Linguistics-based methods typically establish a comprehensive rules to govern the animation process. For example, (Massaro et al. 12) uses dominance functions to map phonemes to facial movements, while JALI (Edwards et al. 2016) factors mouth movements into lip and jaw animation. However, these methods still rely on intermediate representations of phonemes and primarily focus on mouth movements. In contrast, learning-based methods have emerged to address these limitations. For instance, (Taylor et al. 2017) employs deep neural networks to transform phoneme transcriptions into facial animation parameters, proposing a sliding window approach instead of an RNN. Another notable study (Karras et al. 2017) utilizes convolutional neural networks to animate faces directly from audio data.

Recent studies related to our work have concentrated on training neural networks using cross-modal datasets that combine audio and 3D facial meshes. VOCA (Cudeiro et al. 2019), for example, inputs raw audio and speaker style, represented by subject identity, and utilizes temporal convolutions to animate a static mesh template. MeshTalk (Richard et al. 2021) creates a categorical latent space to distinguish between audio-correlated and uncorrelated facial movements. Both tend to overlook long-term audio context due to their reliance on short audio windows. FaceFormer (Fan et al. 2022) addresses this issue by considering long-term audio context with a transformer decoder (Vaswani 2017), enabling it to generate temporally stable animations. In addition, it employs Wav2Vec2.0 (Baevski et al. 2020), a self-supervised pre-trained speech model, which helps mitigate the scarcity of data in existing audio-visual datasets by leveraging large-scale unlabeled speech data to learn rich acoustic and linguistic features. CodeTalker (Xing et al. 2023) integrates a temporal autoregressive model with a latent codebook using VQ-VAE (Van Den Oord, Vinyals et al. 2017), inspired by Learning2Listen (Ng et al. 2022). The above-mentioned methods are deterministic models so their diversity is limited since human speech and facial expressions are variable and dynamic.

### Diffusion Models for Facial Motion Synthesis

Diffusion probabilistic models (Ho, Jain, and Abbeel 2020; Sohl-Dickstein et al. 2015), which differ from previous generative approaches such as GANs (Goodfellow et al. 2020) and VAEs (Kingma 2013), have achieved remarkable results

in various generative tasks. These models adopt a Markov chain to gradually add noise to data samples and subsequently use a neural network to approximate the reverse process to denoise the samples. Given their strong abilities in modeling continuous distributions and one-to-many mapping relationships, diffusion models are particularly suitable for animation tasks.

Diffusion models have been applied in related tasks such as 2D talking face generation (Du et al. 2023). FaceDiffuser (Stan, Haque, and Yumak 2023) was the first to introduce diffusion models to the domain of speech-driven 3D facial animation. DiffPoseTalk (Sun et al. 2024) introduces a speaking style encoder and overcomes the limitations of existing diffusion models that cannot be directly transferred to speech-driven expression animation. GLDiTalker (Lin et al. 2024) proposes a graph-enhanced quantized space and applies a two-stage training strategy. DiffSpeaker (Ma et al. 2024) is one of the most recent and effective works in this field, which introduces a novel biased conditional attention mechanism that utilizes encodings to integrate speaking style and diffusion step information. However, these works typically denoise the entire sequence together, which can lead to performance drops when generating samples that extend beyond the training horizon. In contrast, our model innovates using an autoregressive approach to extract conditions from past motions to guide the diffusion process, enabling flexible generation on arbitrary long sequences.

## Method

Given a speech snippet  $\mathbf{a}^{1:T} = (\mathbf{a}_1, \dots, \mathbf{a}_T)$  and a speaker identity  $s_k$ , our goal is to generate a facial mesh sequence  $\mathbf{x}^{1:T} = (\mathbf{x}_1, \dots, \mathbf{x}_T)$  where each frame  $\mathbf{x}_t \in \mathbb{R}^{V \times 3}$  denotes 3D movements over a template face mesh comprising  $V$  vertices. For simplicity, in the following part, we will denote facial motion by  $\mathbf{x}$  and audio by  $\mathbf{a}$ .

We formulate speech-driven facial animation as a conditional generation problem and propose to use an autoregressive (AR) diffusion model to solve it. An overview of our proposed method is illustrated in Figure 1. We begin by converting facial motion to latent space using VQ-VAE in Section . To overcome the challenge of integrating past motions, we introduce the AR diffusion model, which use the output of the AR condition predictor to guide the diffusion process in Section . In Section , we present the training objectives.

### Learning Latent Space of Facial Motions

To train generative models more easily, it is common to learn a latent space representation of the raw data (Rombach et al. 2022). CodeTalker (Xing et al. 2023) has shown that VQ-VAE can learn a compact and discrete latent space for facial motions. Inspired by this, we adopt a similar VQ-VAE architecture to encode raw facial motion sequences into latent representations.

Our transformer-based VQ-VAE consists of an encoder  $E$ , a decoder  $D$ , and a codebook  $\mathcal{Z} = \{\mathbf{z}_j \in \mathbb{R}^C\}_{j=1}^N$  containing facial motion primitives  $\mathbf{z}$ , where  $N$  represents the size of the codebook. The facial motion data  $\mathbf{x}$  are initially encoded into a continuous feature vector  $\hat{\mathbf{z}} = E(\mathbf{x}) \in$

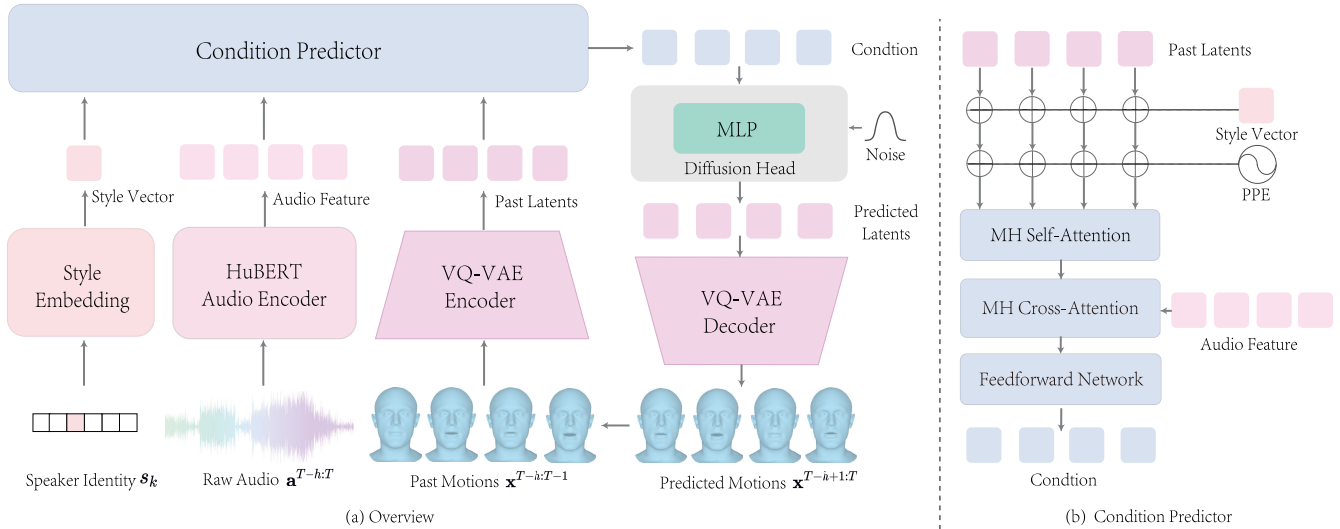


Figure 1: (a) Overview of the pipeline. We employ an AR diffusion model to generate speech-driven 3D facial animations for inputs of arbitrary length. The model first encodes past motions  $\mathbf{x}^{T-h:T-1}$ , raw audio  $\mathbf{a}^{T-h:T}$  and speaker identity  $s_k$  to a dynamic condition. Then the diffusion head leverages this condition to guide the diffusion process. (b) The condition predictor. The condition predictor uses a transformer network to fuse the motion and audio modalities.

$\mathbb{R}^{T \times C}$ . This feature vector is then reshaped into  $\hat{\mathbf{z}}_h \in \mathbb{R}^{T' \times H \times C}$  where  $H$  denotes the number of facial components, and  $T'$  is the number of encoded units, given by  $T' = \frac{T}{H}$ . The quantized feature vector  $\mathbf{z}_q \in \mathbb{R}^{T' \times H \times C}$  is obtained through the quantization operator  $\text{quantize}(\cdot)$ , which calculates the distances between each element of  $\hat{\mathbf{z}}$  and the entries in the codebook, selecting the closest match as follows:

$$\mathbf{z}_q = \text{quantize}(\hat{\mathbf{z}}) = \arg \min_{\mathbf{z}_k \in \mathcal{Z}} \|\hat{\mathbf{z}}_i - \mathbf{z}_k\|_2. \quad (1)$$

Finally, the decoder is employed to reconstruct the facial motions, as defined by Equation 2:

$$\hat{\mathbf{x}} = D(\mathbf{z}_q) = D(\text{quantize}(E(\mathbf{x}))). \quad (2)$$

### Autoregressive Diffusion Model for Facial Animation

In the following sections, we detail how our model generates facial meshes based on a speech snippet  $\mathbf{a}$ , style  $s_k$ , and time step  $t$ . The process begins with input encoding, after which the AR condition predictor estimates the condition for the next frame. Finally, the diffusion head leverages this condition to guide the denoising process. Figure 1 illustrates the entire process.

We first introduce how we encode inputs. Relative studies (Fan et al. 2022; Stan, Haque, and Yumak 2023; Xing et al. 2023) have shown that self-supervised pre-trained speech model features like Wav2Vec2 (Baevski et al. 2020) and HuBERT (Hsu et al. 2021) outperform traditional ones such as MFCC. HuBERT was trained on 960 hours of LibriSpeech (Panayotov et al. 2015) dataset and (Ma et al. 2024) using HuBERT achieves previous state-of-the-art performance. Therefore, we choose HuBERT as the audio encoder in this paper and employ a pre-trained *hubert-base-*

*ls960* version of it. The audio encoder takes the audio snippet  $\mathbf{a}$  as input and outputs a sequence of audio embeddings  $E_a(\mathbf{a}) \in \mathbb{R}^{T \times C}$ . In addition, a simple linear projection layer is used to embed the one-hot style vector  $s_k$  into the latent space as  $E_s(s_k) \in \mathbb{R}^{1 \times C}$ .

**AR Condition Predictor.** Previous work that applies diffusion models typically denoises the entire sequence. However, this approach results in a significant performance drop when generating long sequences that extend beyond the training horizon, especially for long sequences of several thousand frames. To address the challenge of dynamically integrating historical information from past latents and enhancing the model’s generalization capabilities, we propose a novel autoregressive condition predictor. This condition will be used later to guide the diffusion generation process.

During both training and inference, we propose a fixed history length strategy, selecting the start frame and take the next  $h$  frames as the past motion  $\mathbf{x}^{T-h:T-1}$ , which is encoded into latents  $\mathbf{z}_{past}$ . Notably,  $h$  is chosen based on dataset characteristics, ranging from 60 to 120 frames. When the past motion length is less than  $h$  at the beginning of inference, we use the entire sequence:

$$\mathbf{z}_{past} = \text{quantize}(E(\mathbf{x}^{T-h:T-1})) = \mathbf{z}_q^{T-h:T-1}. \quad (3)$$

Then we adopt a transformer decoder to predict the condition  $C_A$  for the next frame  $T$  with a teacher forcing scheme:

$$C_A = \text{TransformerDecoder}(\mathbf{z}_{past}, E_a(\mathbf{a}), E_s(s_k)). \quad (4)$$

where  $\mathbf{z}_{past}$  is the past motion latents,  $E_a(\mathbf{a})$  is the audio embeddings, and  $E_s(s_k)$  is the style embeddings.

As shown in Figure 1 (b), the past motion latents  $\mathbf{z}_{past}$  is first passed through a biased causal multi-head (MH) self-attention layer based on ALiBi (Press, Smith, and Lewis 2021), with a causal mask ensuring that only current or past

information is accessible to prevent information leakage. Then, a multi-head cross-attention layer aligns the motion and audio modalities by combining the outputs of the HuBERT encoder and the MH self-attention layer. A notable design is an alignment mask of the decoder (Ma et al. 2024), which ensures proper alignment of the speech and motion modalities.

**Diffusion Head.** We leverage the condition  $C_A$  from the AR condition predictor to guide the diffusion process in facial animation generation, which is designed to recover latent representations from Gaussian noise based on a conditional distribution  $q(\mathbf{z}_{t-1}|\mathbf{z}_t, C_A)$ . Since this distribution depends on the entire dataset and is intractable, we approximate it using a learnable neural network.

In practice, we train a single-layer MLP to approximate this distribution, ensuring real-time generation speed while maintaining accuracy (Equation 5). Particularly, our model directly predicts the clean sample  $\mathbf{z}_0$ , following MDM (Tevet et al. 2022) and EDGE (Tseng, Castellon, and Liu 2023), as it enables us to offer more precise constraints on facial motions. The denoising process is described as follows:

$$\tilde{\mathbf{z}}_0 = \text{MLP}(\mathbf{z}_t, C_A, E_t(t)), \quad (5)$$

where  $\mathbf{z}_t$  is current noisy latent and  $E_t(t)$  is the time embedding of  $t$ , providing the model with temporal information. After generating the latent  $\tilde{\mathbf{z}}_0$ , we convert it into a facial mesh using the VAE decoder:

$$\tilde{\mathbf{x}} = D(\tilde{\mathbf{z}}_0). \quad (6)$$

## Training Loss and Implementation Details

The training process is divided into two stages: (1) pre-training the VQ-VAE model and (2) training the autoregressive diffusion model.

**Stage 1.** We pre-train the VQ-VAE model following the approach in (Zhang et al. 2018; Razavi, Van den Oord, and Vinyals 2019) for 400 iterations. The encoder and decoder of the VQ-VAE are both a six-layer transformer architecture with eight attention heads, with its feature dimension set to 1024 and the feedforward network dimension set to 1536.

In this stage, we supervise the training with a motion reconstruction loss and a quantization loss similar to CodeTalker (Xing et al. 2023):

$$\mathcal{L}_{stage1} = \lambda_{rec}\mathcal{L}_{rec} + \lambda_{quant}\mathcal{L}_{quant}, \quad (7)$$

where  $\lambda_{rec} = \lambda_{quant} = 1.0$ .

The motion reconstruction loss calculates the L1 distance between the reconstructed facial motion sequence  $\hat{\mathbf{x}}$  and the ground truth facial motion sequence  $\mathbf{x}$ :

$$\mathcal{L}_{rec} = \|\hat{\mathbf{x}} - \mathbf{x}\|_1. \quad (8)$$

The quantization loss contains two intermediate code-level losses that reduce the distance between codebook  $\mathcal{Z}$  and embedded features  $\hat{\mathbf{z}}$ :

$$\mathcal{L}_{quant} = \|sg(\hat{\mathbf{z}}) - \mathbf{z}_q\|_2^2 + \beta\|\hat{\mathbf{z}} - sg(\mathbf{z}_q)\|_2^2, \quad (9)$$

where  $\beta = 0.25$  is a weighting hyperparameter controlling the update speed of codebook and encoder, and  $sg(\cdot)$  stands

for a stop-gradient operation which is defined as identity with zero partial derivatives in the backward propagation. Note that the quantization process (Equation 1) is not differentiable, we employ the straight-through gradient estimator following (Bengio, Léonard, and Courville 2013), to copy the gradients from the decoder’s input to the encoder’s output.

**Stage 2.** Next, we train the AR condition predictor and diffusion head jointly for 200 iterations while keeping the VQ-VAE motion encoder fixed. The AR condition predictor is a two-layer transformer-based decoder with four attention heads, while the diffusion head is a single-layer MLP sharing the same hidden size. For this phase, we adopt a scaled linear schedule (Ho, Jain, and Abbeel 2020) with a total of  $N = 1000$  steps, and we employ the Denoising Diffusion Implicit Models (DDIM) (Song, Meng, and Ermon 2020), using only 50 steps to sample during inference.

To better improve the quality of latents, we introduce a new latent loss function inspired by MAR (Li et al. 2024) that calculates the L1 distance between the generated latents  $\tilde{\mathbf{z}}$  and the ground truth latents  $\mathbf{z}_q$ :

$$\mathcal{L}_{latent} = \|\tilde{\mathbf{z}}_0^{T-h+1:T} - \mathbf{z}_q^{T-h+1:T}\|_1. \quad (10)$$

We then apply the following geometric losses in the 3D space: the vertex loss  $\mathcal{L}_{vert}$  (Cudeiro et al. 2019) for the positions of the mesh vertices and the velocity loss  $\mathcal{L}_{vel}$  (Cudeiro et al. 2019) for better temporal consistency:

$$\mathcal{L}_{vert} = \|\tilde{\mathbf{x}}^{T-h+1:T} - \mathbf{x}^{T-h+1:T}\|_2^2, \quad (11)$$

$$\mathcal{L}_{vel} = \left\| \left( \tilde{\mathbf{x}}^{T-h+1:T} - \tilde{\mathbf{x}}^{T-h:T-1} \right) - \left( \mathbf{x}^{T-h+1:T} - \mathbf{x}^{T-h:T-1} \right) \right\|_2^2. \quad (12)$$

In summary, we combine all the aforementioned losses and apply a weighted sum to obtain the final loss function for stage 2 as:

$$\mathcal{L}_{stage2} = \lambda_{latent}\mathcal{L}_{latent} + \lambda_{vert}\mathcal{L}_{vert} + \lambda_{vel}\mathcal{L}_{vel}, \quad (13)$$

where  $\lambda_{latent} = \lambda_{vert} = \lambda_{vel} = 1.0$ . All models are trained using the AdamW (Loshchilov, Hutter et al. 2017) optimizer with a batch size of 1 and a learning rate of 0.0001 on a single Nvidia RTX 4090 GPU. More details of the network architecture are provided in the Appendix.

## Experiments and Results

### Datasets

We use two publicly available datasets, BIWI (Fanelli et al. 2010) and VOCASET (Cudeiro et al. 2019), to evaluate the effectiveness our method. Both datasets provide 3D scan-audio pairs of utterances spoken in English.

**BIWI Dataset.** The BIWI dataset is a comprehensive collection of affective speech and corresponding detailed 3D facial geometries. It consists of 14 participants who read 40 English sentences, each recorded in neutral and emotional contexts. The 3D facial data is captured at a rate of 25 fps, with each frame containing 23,370 vertices. The average duration of each sequence is approximately 4.67 seconds.

**VOCASET Dataset.** The VOCASET dataset includes 480 paired audio-visual sequences recorded from 12 subjects. The facial motion is captured at 60 frames per second and typically lasts about 4 seconds. Unlike the BIWI dataset, each 3D face mesh in VOCASET is registered to the FLAME (Li et al. 2017) topology, which consists of 5023 vertices.

### Evaluation Metrics

Following previous works (Fan et al. 2022; Xing et al. 2023), we adopt two metrics for quantitative evaluation.

**Lip Vertex Error (LVE).** We measure lip synchronization following (Richard et al. 2021) by calculating the L2 error between the predicted and ground-truth lip vertices.

**Face Dynamics Distance (FDD).** We follow (Xing et al. 2023) evaluating the deviation of upper face motions by calculating each upper face vertex’s standard deviation of the element-wise L2 norm along the temporal axis. Smaller FDD indicates that the predicted expressions exhibit high consistency with the natural trends of facial dynamics.

**Mouth Open Difference (MOD).** Following (Sun et al. 2024), we compute the average absolute difference in mouth opening between the predicted and ground truth. A smaller value indicates better alignment in the degree of mouth opening. This helps reflect whether the model generates over-smooth mouth movements.

### Quantitative Evaluation

**Baselines.** We compare our method with the following state-of-the-art methods: VOCA (Cudeiro et al. 2019), MeshTalk (Richard et al. 2021), FaceFormer (Fan et al. 2022), CodeTalker (Ng et al. 2022), FaceDiffuser (Stan, Haque, and Yumak 2023), DiffSpeaker (Ma et al. 2024), Imitator (Thambiraja et al. 2023) and DiffPoseTalk (Sun et al. 2024). Among these, VOCA uses a MLP as decoder while MeshTalk uses a U-Net, both of which use convolutional neural networks to extract features from audio. FaceFormer, CodeTalker and Imitator use transformers to model the long-term context. FaceDiffuser uses diffusion to generate facial animation and applies GRU as the denoiser. DiffSpeaker chooses transformer decoder as the backbone network and outperforms previous methods. DiffPoseTalk predicts FLAME parameters directly. To enable comparison under our vertex-based setting, we adapt its official implementation to output mesh vertices. For all methods, we follow their original training and testing protocols as closely as possible, and apply official pre-trained weights if available. For speakers unseen during training (i.e., not learned in the style embedding space), we generate multiple animations by conditioning on each speaker identity in training set separately, then compute evaluation metrics for each and report the average as the final result.

We apply all metrics to the BIWI dataset but only LVE and MOD for the VOCASET dataset based on its limited facial expression variation. As shown in Table 1, our model achieves the best results on both LVE and FDD, demonstrating its ability to produce accurate lip movements and

Methods	VOCASET		BIWI		
	LVE ↓	MOD ↓	LVE ↓	FDD ↓	MOD ↓
VOCA	4.9245	4.7131	6.5563	8.1816	33.7083
MeshTalk	4.5441	4.4520	5.9181	5.1025	18.9012
FaceFormer	4.1090	4.2001	5.3077	4.6408	16.7510
CodeTalker	3.9445	4.1834	4.7914	4.1170	13.4410
FaceDiffuser	4.1089	3.8069	4.2986	3.9098	8.6695
DiffSpeaker	3.1478	3.5339	4.2829	3.8535	<b>8.4091</b>
Imitator	3.1352	3.6610	4.3515	3.8066	10.3524
DiffPoseTalk	3.9110	3.8819	5.3128	3.8256	12.6930
Ours	<b>2.7206</b>	<b>3.4987</b>	<b>4.2504</b>	<b>3.6690</b>	8.5208

Table 1: Quantitative comparison of StreamingTalker with baseline methods on two benchmark datasets. For VOCASET, LVE and MOD values are in units of  $10^{-5}$ mm. For BIWI, LVE and MOD values are in units of  $10^{-4}$ mm, while FDD values are in units of  $10^{-5}$ mm. Among them, bold indicates the best performance. ↓ means lower is better.

Methods	LVE ↓	FDD ↓	MOD ↓
FaceFormer	5.4079	5.1418	20.7980
CodeTalker	5.7064	6.6470	15.1021
FaceDiffuser	5.3125	4.2674	9.0024
DiffSpeaker	5.2213	4.7980	8.9528
Ours	<b>4.4596</b>	<b>3.8912</b>	<b>8.8017</b>

Table 2: Comparison of results on long sequences concatenated from BIWI-Test-B. The generated facial motion lengths are set to 2000 frames. The units used here are same as in Table 1.

natural upper-face dynamics synchronized with speech. Although the MOD is slightly higher than that of DiffSpeaker on short sequences, the difference is marginal and remains within an acceptable perceptual range.

Additionally, we evaluate the model’s ability to generate extended facial motion sequences far beyond the training horizon. Specifically, we use BIWI sequences of 2000 frames (i.e., over 60 seconds) to test long-term generalization. As shown in Table 2, our model consistently outperforms previous methods across all metrics, including MOD, indicating better temporal stability and coherence over long durations. This result highlights the robustness of our autoregressive framework in modeling long-term dependencies without sacrificing motion quality or expressiveness.

### Qualitative Evaluation

In the qualitative evaluation, we visually compare our method with FaceFormer, CodeTalker and DiffSpeaker. As shown in Figure 2, the animation results of different methods reveal that our approach generates lip movements that are more accurate and better synchronized with the audio. The mouth opens and closes more naturally compared to our competitors. For example, our method excels in handling rounded vowels, as demonstrated in the word “body” and “now”. Besides, for words like “quick”, “chaps” and “just”




Figure 2: Qualitative comparison with the state-of-the-arts. The left side shows results on the VOCASET-Test dataset, while the right side shows results on the BIWI-Test-B dataset. Red words indicate phonemes being pronounced. Compared to other methods, our approach produces more natural lip shapes, with rounder mouth formations when pronouncing vowels like 'a', 'o', and 'u', and better lip closure for bilabial consonants such as 'm' and 'p'.

the model accurately generates a pursed lip shape. Furthermore, our method ensures proper lip closure in words like "ambiguous" and "parents", accurately modeling the /m/ and /p/ sounds that require the lips to be fully closed. For more visual comparisons, please refer to our supplementary video.

### Ablation Study

As shown in Table 3, we compare our method with five main ablation settings:

(1) **w/o Diffusion Head.** To validate the effectiveness of the diffusion process, we remove the diffusion head and directly predict facial motion from the condition using an MLP. Without the diffusion head, the model's ability to capture the complex distribution of facial motion is reduced, resulting in degraded performance across all metrics.

(2) **Effect of Motion Latent Encoding.** To evaluate the impact of motion latent encoding, we first remove the VQ-VAE and directly denoise the motion. This leads to a significant drop in motion quality, highlighting the importance of latent representation. We further replace the VQ-VAE with a standard VAE to assess the model's robustness to the choice of encoder. The results show only marginal performance differences, suggesting that our model remains stable even with

continuous latent representations.

(3) **Use All History Motions.** To evaluate the effectiveness of our fixed-length history strategy, we instead use all previous frames as historical context. This leads to slightly worse performance. We attribute this to a mismatch between training and testing conditions: the BIWI training sequences are relatively short, while test sequences are much longer. Using all history at test time exposes the model to significantly longer temporal dependencies than seen during training, causing a distribution shift that harms performance. In contrast, our fixed-window approach enforces a consistent context length, leading to better generalization across varying sequence durations.

(4) **w/o Cross-Attention.** To evaluate the role of cross-attention layers, we remove all of them, eliminating the interaction between audio and motion features. Without audio guidance, the generated facial meshes become nonsensical and lack semantic consistency.

(5) **w/o Self-Attention.** To examine the impact of self-attention, we remove all self-attention layers. This results in a noticeable decline in global consistency, particularly in the upper face, where movements lose coherence, leading to lower FDD scores.

Additional qualitative results for the ablation study can be

Methods	BIWI			VOCASET	
	LVE ↓ ( $\times 10^{-4}$ mm)	FDD ↓ ( $\times 10^{-5}$ mm)	MOD ↓ ( $\times 10^{-5}$ mm)	LVE ↓ ( $\times 10^{-5}$ mm)	MOD ↓ ( $\times 10^{-5}$ mm)
w/o Diffusion Head	4.4297	5.3934	12.4698	3.5070	3.6011
w/o VQ-VAE	8.4578	9.7510	9.2511	4.9521	9.7510
Use VAE as encoder	4.2686	3.7010	8.4292	2.7516	5.2660
Use all history motions	4.7598	4.0141	8.7296	2.8168	3.5637
w/o Cross-Attention	11.7510	6.9653	21.2168	7.4660	10.6860
w/o Self-Attention	4.3025	4.1083	15.5354	3.0321	3.7499
<b>Ours</b>	<b>4.2504</b>	<b>3.6690</b>	<b>8.5208</b>	<b>2.7206</b>	<b>3.4987</b>

Table 3: Ablation study. We compare our method with six variants to validate our main design choices. Among them, bold indicates the best results. ↓ means lower is better. Mean values are reported.

found in our supplementary materials.

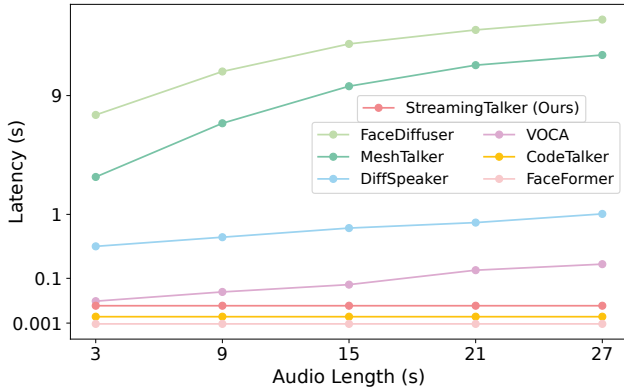


Figure 3: Inference latency for 3-27 second audio clips. The figure compares the performance of various models, including full sequence diffusion models (DiffSpeaker, FaceDiffuser), deterministic models (VOCA, MeshTalk), and AR models (FaceFormer, CodeTalker). Our model outperforms all non-AR models in terms of inference speed, maintaining consistent latency regardless of audio length.

### Real-time Application

Inference latency is a critical factor in real-time applications, as it directly affects the user experience. In this paper, we define inference latency as the time elapsed from the input of audio to the generation of the first frame of animation output. We measured the inference latency for audio inputs of varying lengths on a 4090 GPU. The results, as shown in Figure 3, demonstrate that our method significantly outperforms all full sequence diffusion methods in terms of latency. Additionally, our method maintains consistent inference latency across varying audio lengths. This improvement is attributed to the unique design of our AR diffusion model, which supports streaming input. Unlike traditional diffusion methods that process the entire sequence at once, our model can begin rendering the output as soon as the first frame is processed. This process takes only 25 ms, thus meeting the requirements for real-time rendering.

We further build a real-time demo leveraging the streaming capability of our model. In this demo, a user speaks to ask a question, and the system first generates a response using a large language model (LLM). The response is then converted to speech via Google TTS (pndurette 2021), which is fed into our AR diffusion model to generate facial animation. Thanks to the model’s low-latency design, audio is processed in a streaming manner, with animation frames generated and rendered simultaneously at up to 40 FPS.

Further details on per-component latency, implementation of our real-time demo, and the video demonstration can be found in the supplementary material.

### Discussions

Although our proposed approach achieves significant improvements in generating 3D facial animations from speech, it still has some limitations. Firstly, the current dataset consists of relatively short sequences, which are not sufficient to capture the full range of facial expressions a person can exhibit. As a result, our model is unable to extract long-term contextual information, so a larger and more diverse dataset could potentially improve the model’s performance. Secondly, our model currently lacks emotional information, which could be a valuable addition to future research. Incorporating emotional cues, such as EMOCA (Daněček, Black, and Bolkart 2022), could enhance the model’s ability to generate more expressive animations and further improve the performance.

### Conclusion

In this work, we introduce an AR diffusion model for speech-driven 3D facial animation. Our model dynamically integrates historical motion data, enhancing adaptability and context awareness compared to full sequence diffusion methods. This approach mitigates performance drops when generating sequences beyond the training horizon. Quantitative results show our method outperforms state-of-the-art techniques. We also developed an interactive real-time demo integrating an LLM for customizable attribute control and supporting streaming input, ensuring smooth, real-time rendering. Extensive ablation studies validate our network architecture and design choices.

## Ethical Statement

As our approach enables realistic synthesis of facial motion from audio, we acknowledge potential risks of misuse, such as in generating deceptive or unauthorized content. We emphasize that our work is intended for legitimate applications, including virtual communication, accessibility tools, and creative production.

## Acknowledgments

This work was partially supported by the following grants: National Key R & D Program of China (No. 2024YFB2809102), Zhejiang Provincial Natural Science Foundation of China (No. LR25F020003), Information Technology Center and State Key Lab of CAD & CG, Zhejiang University.

## 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.
- Bengio, Y.; Léonard, N.; and Courville, A. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*.
- Cudeiro, D.; Bolkart, T.; Laidlaw, C.; Ranjan, A.; and Black, M. J. 2019. Capture, learning, and synthesis of 3D speaking styles. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10101–10111.
- Daněček, R.; Black, M. J.; and Bolkart, T. 2022. Emoca: Emotion driven monocular face capture and animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20311–20322.
- Du, C.; Chen, Q.; He, T.; Tan, X.; Chen, X.; Yu, K.; Zhao, S.; and Bian, J. 2023. Dae-talker: High fidelity speech-driven talking face generation with diffusion autoencoder. In *Proceedings of the 31st ACM International Conference on Multimedia*, 4281–4289.
- Edwards, P.; Landreth, C.; Fiume, E.; and Singh, K. 2016. Jali: an animator-centric viseme model for expressive lip synchronization. *ACM Transactions on graphics (TOG)*, 35(4): 1–11.
- Fan, Y.; Lin, Z.; Saito, J.; Wang, W.; and Komura, T. 2022. Faceformer: Speech-driven 3d facial animation with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18770–18780.
- Fanelli, G.; Gall, J.; Romsdorfer, H.; Weise, T.; and Van Gool, L. 2010. A 3-d audio-visual corpus of affective communication. *IEEE Transactions on Multimedia*, 12(6): 591–598.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Hsu, W.-N.; Bolte, B.; Tsai, Y.-H. H.; Lakhotia, K.; Salakhutdinov, R.; and Mohamed, A. 2021. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM transactions on audio, speech, and language processing*, 29: 3451–3460.
- 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. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Li, T.; Bolkart, T.; Black, M. J.; Li, H.; and Romero, J. 2017. Learning a model of facial shape and expression from 4D scans. *ACM Trans. Graph.*, 36(6): 194–1.
- Li, T.; Tian, Y.; Li, H.; Deng, M.; and He, K. 2024. Autoregressive Image Generation without Vector Quantization. *arXiv preprint arXiv:2406.11838*.
- Lin, Y.; Fan, Z.; Xiong, L.; Peng, L.; Li, X.; Kang, W.; Wu, X.; Lei, S.; and Xu, H. 2024. GLDiTalker: Speech-Driven 3D Facial Animation with Graph Latent Diffusion Transformer. *arXiv preprint arXiv:2408.01826*.
- Loshchilov, I.; Hutter, F.; et al. 2017. Fixing weight decay regularization in adam. *arXiv preprint arXiv:1711.05101*, 5.
- Ma, Z.; Zhu, X.; Qi, G.; Qian, C.; Zhang, Z.; and Lei, Z. 2024. DiffSpeaker: Speech-Driven 3D Facial Animation with Diffusion Transformer. *arXiv preprint arXiv:2402.05712*.
- Massaro, D.; Cohen, M.; Tabain, M.; Beskow, J.; and Clark, R. 2012. 12 Animated speech: research progress and applications.
- Ng, E.; Joo, H.; Hu, L.; Li, H.; Darrell, T.; Kanazawa, A.; and Ginosar, S. 2022. Learning to listen: Modeling non-deterministic dyadic facial motion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20395–20405.
- Panayotov, V.; Chen, G.; Povey, D.; and Khudanpur, S. 2015. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 5206–5210. IEEE.
- pndurette. 2021. gTTS: Google Text-to-Speech. Accessed: 2025-01-23.
- Press, O.; Smith, N. A.; and Lewis, M. 2021. Train short, test long: Attention with linear biases enables input length extrapolation. *arXiv preprint arXiv:2108.12409*.
- Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; and Sutskever, I. 2021. Zero-shot text-to-image generation. In *International conference on machine learning*, 8821–8831. Pmlr.
- Razavi, A.; Van den Oord, A.; and Vinyals, O. 2019. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32.
- Richard, A.; Zollhöfer, M.; Wen, Y.; De la Torre, F.; and Sheikh, Y. 2021. Meshtalk: 3d face animation from speech using cross-modality disentanglement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1173–1182.

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.

Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, 2256–2265. PMLR.

Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.

Stan, S.; Haque, K. I.; and Yumak, Z. 2023. Facediffuser: Speech-driven 3d facial animation synthesis using diffusion. In *Proceedings of the 16th ACM SIGGRAPH Conference on Motion, Interaction and Games*, 1–11.

Sun, Z.; Lv, T.; Ye, S.; Lin, M.; Sheng, J.; Wen, Y.-H.; Yu, M.; and Liu, Y.-j. 2024. Diffposetalk: Speech-driven stylistic 3d facial animation and head pose generation via diffusion models. *ACM Transactions on Graphics (TOG)*, 43(4): 1–9.

Taylor, S.; Kim, T.; Yue, Y.; Mahler, M.; Krahe, J.; Rodriguez, A. G.; Hodgins, J.; and Matthews, I. 2017. A deep learning approach for generalized speech animation. *ACM Transactions On Graphics (TOG)*, 36(4): 1–11.

Tevet, G.; Raab, S.; Gordon, B.; Shafir, Y.; Cohen-Or, D.; and Bermano, A. 2022. Human motion diffusion model. *arXiv 2022. arXiv preprint arXiv:2209.14916*.

Thambiraja, B.; Habibie, I.; Aliakbarian, S.; Cosker, D.; Theobalt, C.; and Thies, J. 2023. Imitator: Personalized speech-driven 3d facial animation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 20621–20631.

Tseng, J.; Castellon, R.; and Liu, K. 2023. Edge: Editable dance generation from music. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 448–458.

Van Den Oord, A.; Vinyals, O.; et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.

Vaswani, A. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.

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.

Zhang, H.; Starke, S.; Komura, T.; and Saito, J. 2018. Mode-adaptive neural networks for quadruped motion control. *ACM Transactions on Graphics (TOG)*, 37(4): 1–11.