

FaceTalk: Audio-Driven Motion Diffusion for Neural Parametric Head Models

Shivangi Aneja¹ Justus Thies^{2,3} Angela Dai¹ Matthias Nießner¹

¹Technical University of Munich ²MPI-IS, Tübingen ³TU Darmstadt




Figure 1. Given an input speech signal, we propose a diffusion-based approach to synthesize high-quality and temporally consistent 3D motion sequences of high-fidelity human heads as neural parametric head models. Our method can generate a diverse set of expressions (including wrinkles and eye blinks) and the generated mouth motion is temporally synchronized with the given audio signal.

Abstract

We introduce FaceTalk¹, a novel generative approach designed for synthesizing high-fidelity 3D motion sequences of talking human heads from input audio signal. To capture the expressive, detailed nature of human heads, including hair, ears, and finer-scale eye movements, we propose to couple speech signal with the latent space of neural parametric head models to create high-fidelity, temporally coherent motion sequences. We propose a new latent diffusion model for this task, operating in the expression space of neural parametric head models, to synthesize audio-driven realistic head sequences. In the absence of a dataset with corresponding NPHM expressions to audio, we optimize for these correspondences to produce a dataset of temporally-optimized NPHM expressions fit to audio-video recordings of people talking. To the best of our knowledge, this is the first work to propose a generative approach for realistic and high-quality motion synthesis of volumetric human heads, representing a significant advancement in the field of audio-driven 3D animation. Notably, our approach stands out in its ability to generate plausible motion sequences that can produce high-fidelity head animation coupled with the

NPHM shape space. Our experimental results substantiate the effectiveness of FaceTalk, consistently achieving superior and visually natural motion, encompassing diverse facial expressions and styles, outperforming existing methods by 75% in perceptual user study evaluation.

1. Introduction

Modeling 3D animation of humans has a wide range of applications in the realm of digital media, including animated movies, computer games, and virtual assistants. In recent years, there have been numerous works proposing generative approaches for motion synthesis of human bodies, enabling the animation of human skeletons conditioned on various signals such as action [3, 25, 46], language [1, 4, 36, 47, 64, 77] and music [2, 67]. While human faces are critical to synthesis of humans, generative synthesis of 3D faces in motion has focused on 3D morphable models (3DMMs) leveraging linear blendshapes [7, 40] to represent head motion and expression. Such models characterize a disentangled space of head shape and motion, but lack the capacity to comprehensively represent the complexity and fine-grained details of human face geometry in motion (e.g., hair, skin furrowing during motion, etc.).

Thus, we propose to represent animated head sequences

¹Project Page: <https://shivangi-aneja.github.io/projects/faceTalk>

with a volumetric 3D head representation, leveraging the expressive representation space of neural parametric head models (NPHMs) [23, 24]. NPHMs offer a flexible representation capable of handling complex and irregular facial expressions (e.g., blinking, skin creasing), along with a high-fidelity shape space including the head, hair, and ears, making them a much more suitable choice for face animation. We address the challenging task of creating an audio-conditional generative animation model for this volumetric representation.

We design the first transformer-based latent diffusion model for audio-driven head animation synthesis. Our diffusion model operates in the latent NPHM expression space to generate temporally coherent expressions consistent with an input audio signal represented with Wave2Vec 2.0 [6]. In the absence of paired audio-NPHM data, we optimize for corresponding NPHM expressions to fit to multi-view video recordings of people speaking, generating train supervision for our task. As NPHMs are designed for static (frame-by-frame) expressions without temporal consistency, we employ both geometric and temporal priors to produce temporally consistent optimized motion sequences. This enables training our audio-head diffusion model to synthesize realistic speech-conditioned 3D head sequences, capable of capturing high-frequency details like wrinkles and eye blinks present in the face region. Our method takes first step towards simplifying the task of high-fidelity facial motion generation of 3D faces for content creation applications.

In summary, we propose the first latent diffusion model for the creation of audio-conditioned animations of volumetric avatars. By producing volumetric head animation, our generative model is highly expressive yet efficient compared to existing 3D methods. We also demonstrate control over motion style, using classifier-free guidance to adjust the strength of the stylistic expression.

2. Related Work

Facial Animation. Our proposed method is the first work for audio-conditioned latent diffusion of volumetric head avatars. There is a large corpus of research works in the field of 2D audio-driven facial animation operating on RGB videos, synthesizing 2D sequences directly [10–13, 16, 27, 33, 34, 48, 56, 59, 61, 70, 71, 73, 79, 80]. However, these methods operate in pixel space and do not produce any geometric information. Another line of work also operating on RGB videos but using intermediate 3D representations are based on 3DMMs [19, 32, 57, 62, 66]. Although these methods generate RGB videos in the end, they use 3DMMs which produces very sparse geometric information.

Recent works based on radiance fields [22, 26, 37, 41, 55, 76] have also gained popularity due to their capability to model densities directly from images. These methods gen-

erate impressive RGB videos but the underlying geometry learned is highly imperfect (no segregation between background and facial region) as shown by Chan *et al.* [9]. Additionally, these methods require identity-specific training, and thus can not be used for content creation applications.

Learning animation of 3D meshes directly is much more promising, but only a handful of methods exist [14, 15, 21, 35, 44, 51, 65, 75]. A vast majority of these works model speech-conditioned animation for an artist-designed template mesh. Although these methods can match the facial motion with the speech signal, one limitation of these methods is their incapability to represent fine-scale details present in faces. Another downside of these methods is that these approaches learn a deterministic model producing no/muted motion in the upper region of the face, thus limiting them from being able to produce realistic motion. In this work, we solve these issues by proposing a generative model that can operate in the compact and detailed latent space of neural parametric head models, thus capable of representing fine-scale facial details and synthesizing a diverse set of expressions and speaking styles.

Diffusion Models for Generative Synthesis. In recent years, diffusion models have experienced a surge of interest as highly expressive and efficient generative models. These models have demonstrated strong performances as a generative model for a variety of domains such as images [17, 29, 49, 52, 53, 78], videos [8, 18, 20, 69], speech [30, 38, 39] and motion [2, 5, 36, 64, 67, 81].

2D facial animation has also seen some progress with diffusion models [56, 59]. DiffusedHeads [59] operates in pixel space, making the sampling process very slow. Although DiffTalk [56] learns the diffusion process in the latent space, the sampling is performed in an autoregressive fashion, limiting sampling efficiency. Moreover, these methods operate on RGB videos, and though they achieve high-quality results on a per-frame basis, consistency and coherence across different timesteps remain challenging due to temporal jitter. Hence, these methods can not be applied directly to the parametric head models. Concurrent to our work, DiffPoseTalk [60] and FaceDiffuser [58] propose a diffusion-based approach for animating 3D meshes from speech signal. However, these methods require additional conditions like style embedding, generate sequences in an autoregressive fashion coupled with diffusion denoising making them very slow and still they can not produce fine-scale facial details due to the use of 3DMMs which are limited in expressivity. Concurrent to ours, DPHM [63] also utilizes diffusion to regularize NPHM expressions from depth, operating on a frame-by-frame basis. In contrast, we synthesize the entire sequence simultaneously making it much faster and operate directly in the latent space of highly expressive NPHMs, producing high-fidelity and temporally consistent results.

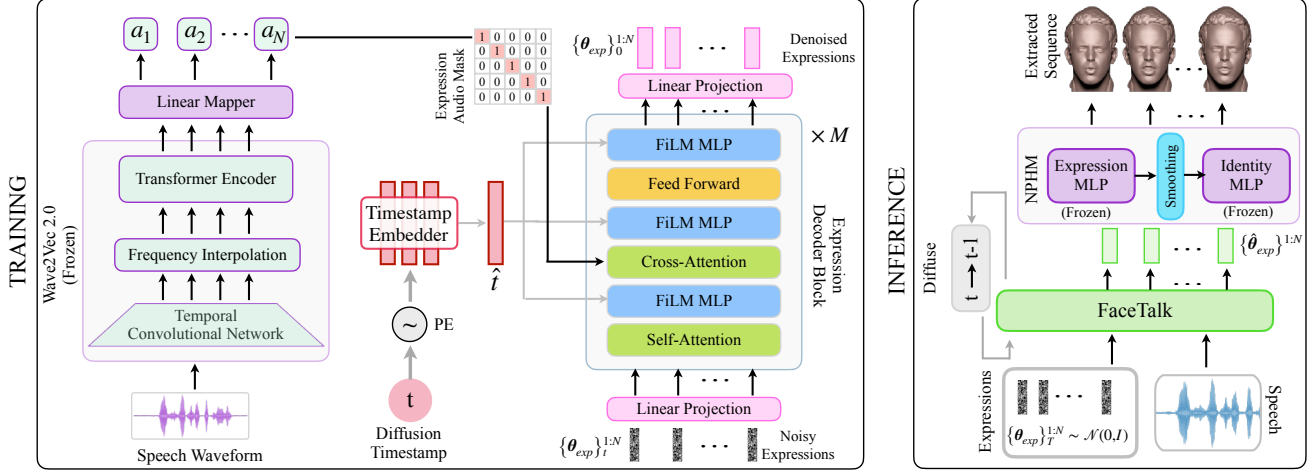


Figure 2. **Pipeline Overview.** FaceTalk uses frozen Wave2Vec 2.0 [6] to extract audio embeddings from a speech signal. The diffusion timestamp is embedded using a timestamp embedder. The expression decoder employs a multi-head transformer decoder [68] with FiLM [45] layers, interleaved between Self-Attention, Cross Attention, and FeedForward layers, to incorporate diffusion timestamp. During training, the model is trained to denoise the noisy expression sequences from timestamp t . At inference, FaceTalk denoises the gaussian noise sequence $\{\theta_{exp}\}_T^{1:N} \sim \mathcal{N}(0, \mathbf{I})$ iteratively until $t = 0$, yielding the estimated final sequence $\{\hat{\theta}_{exp}\}_T^{1:N}$. These are then input to the frozen NPHM model, utilizing facial smoothing, and mesh sequences are extracted using MC [42].

3. Preliminaries

NPHM. In contrast to traditional 3D morphable models, which remain limited in expressive detail, we employ the Neural Parametric Head Model (NPHM) representation [23, 24]. Similar to traditional 3DMMs, NPHM disentangles a human face into an identity and expression space; however, the shape and expression space are volumetric, enabling expressive capture of details such as fine-scale eye movements, hair, ears, etc. NPHM uses two auto-decoder style neural networks (a) Identity Network $\{\mathcal{I}\}$ to represent the overall facial shape and (b) Expression Network $\{\mathcal{E}\}$ to represent the facial movements such as jaw pose, wrinkles, eyeblinks, etc., jointly denoted as $\mathcal{F} = \{\mathcal{I}, \mathcal{E}\}$. The identity and expressions latent codes $\{\theta_{id}, \theta_{exp}\}$ define the facial shape and expression respectively. Unlike 3DMMs which predict a fixed-topology mesh that can not model fine surface details, NPHM can handle different hairstyles, wrinkles and complex facial expressions. NPHM represents identities with a signed distance field (SDF) in canonical space and models the expressions as deformations onto the facial region. Mathematically, it can be defined as:

$$\mathcal{F}(x_i, \theta_{id}, \theta_{exp}) \rightarrow s_i : \mathbb{R}^3 \times \mathbb{R}^{\theta_{id}} \times \mathbb{R}^{\theta_{exp}} \rightarrow \mathbb{R}, \quad (1)$$

where $x_i \in \mathbb{R}^3$ represents the query points, $\theta_{id} \in \mathbb{R}^{1344}$ represents identity latent code, $\theta_{exp} \in \mathbb{R}^{200}$ represents expression latent code and s_i denotes the predicted SDF. The mesh is then extracted as the zero-level iso-surface decision boundary using Marching Cubes [42].

Diffusion. Diffusion models are a class of generative mod-

els that consist of a forward and reverse process. The forward process converts the original structured data distribution into Gaussian noise, modeled following a fixed Markov chain as:

$$q(x_{1:T} | x_0) = \prod_{t=1}^T q(x_t | x_{t-1}), \quad (2)$$

where $q(x_t | x_{t-1})$ denotes the forward process adding white noise to the original data distribution x_0 . Mathematically, the forward process can be written as:

$$q(x_t | x_0) \sim \mathcal{N}(\sqrt{\bar{\alpha}_t} x_0, (1 - \bar{\alpha}_t) \mathbf{I}), \quad (3)$$

where $\bar{\alpha}_t \in (0, 1)$ are constants following a decreasing cosine schedule such that when $\bar{\alpha}_t$ tends to 0, we can approximate $x_T \sim \mathcal{N}(0, \mathbf{I})$. The reverse diffusion process progressively denoises samples from $x_T \sim \mathcal{N}(0, \mathbf{I})$ into samples from a learned distribution x_0 . In our experiments, we learn a generative model \mathcal{G}_θ to reverse the forward diffusion process by learning to estimate $\mathcal{G}_\theta(x_t, t, c) = \hat{x} \approx x$, where x refers to predicting the cleaned samples directly, x_t refers to noisy input at diffusion timestamp t . Given a conditional signal c , we optimize model parameters θ for all diffusion timestamps t using the following objective:

$$\mathcal{L}_\theta = \mathbb{E}_{x,t} [\|x - \mathcal{G}_\theta(x_t, t, c)\|_2^2]. \quad (4)$$

4. Method

FaceTalk performs high-fidelity and temporally consistent generative synthesis of motion sequences of heads, condi-

tioned on audio signal. In order to characterize complex face motions and fine-scale movements, we synthesize realistic heads in the latent expression space of a parametric model for volumetric head representations, i.e., neural parametric head models (NPHMs) [23]. We, thus, develop a speech-conditioned latent diffusion model to synthesize temporally coherent head expression sequences that are coupled with the NPHM shape space to produce complex, realistic head animations of different identities. An overview of our approach is illustrated in Fig. 2.

Audio Encoding. We employ a state-of-the-art pre-trained speech model Wave2Vec 2.0 [6] to encode the audio signal. Specifically, we first use the audio feature extractor made up of temporal convolution layers (TCN) to extract audio feature vectors $\{a_i\}_{i=1}^{N_a}$ from the raw waveform. This is followed by the Frequency Interpolation layer that aligns the input audio signal $\{a_i\}_{i=1}^{N_a}$ (captured at frequency $f_a = 16\text{kHz}$) with our dataset $\{a_i\}_{i=1}^{N_e}$ (captures at framerate $f_e = 24\text{Hz}$). Finally, a stacked multi-layer transformer encoder network processes these resampled features and outputs the processed and aligned audio feature vectors. The audio encoder is initialized with the pre-trained wav2vec 2.0 weights, followed by a feedforward layer to project the aligned audio features into the latent space of our expression decoder model. These aligned audio features are denoted as:

$$\mathbf{A}^{1:N} = \{a_i\}_{i=1}^N. \quad (5)$$

These audio features $\mathbf{A}^{1:N}$ are then fed to the cross-attention layers of the expression decoder via the expression-audio alignment mask \mathcal{M} to learn speech-conditioned expression features during training.

Expression Encoding. We train our expression decoder network on optimized NPHM expression sequences $\{\theta_{exp}\}^{1:N}$ (obtained in Section 5), where N refers to number of frames in a sequence. We train a diffusion-based stacked multi-layer transformer [68] decoder network to synthesize facial expressions in the latent space of the NPHM model. During training, following forward diffusion (Eq. 3) we add noise for a randomly sampled diffusion timestamp $t \sim \text{Uniform}(0, T)$ to create noisy expression codes $\{\theta_{exp}\}_t^{1:N}$. These noisy expression codes are then projected to the latent space of our model via a linear layer, followed by a stack of transformer decoder blocks, and then projected back to the original NPHM space using another linear layer. To embed diffusion timestamp in the latent space of our model, we apply sinusoidal embedding and process it through a three-layer MLP. Next, to fuse diffusion timestamp into the model, we use a one-layer FiLM (feature-wise linear modulation) network [45] between the multi-head self-attention, multi-head cross-attention and feedforward layers of the transformer decoder block, which

is critical for the model to produce high-quality expressions as we show in the results. We leverage a look-ahead binary target mask $\mathcal{T} \in \mathbb{R}^{N \times N}$ in the multi-head self-attention layers to prevent the model from peeking into the future expression codes. Mathematically, it can be written as:

$$\mathcal{T}_{ij} = \begin{cases} True & \text{if } i \leq j \\ False & \text{else} \end{cases} \quad (6)$$

where \mathcal{T}_{ij} refers to the $(i, j)^{th}$ element of the matrix and $1 \leq i, j \leq N$. The audio features $\mathbf{A}^{1:N}$ are fused into the network via the multi-head cross-attention layers. We leverage the expression-audio alignment mask \mathcal{M} to fuse the audio features into the network to ensure that mouth poses learned by the expression codes are consistent with the speech signal. The binary mask $\mathcal{M} \in \mathbb{R}^{N \times N}$ is Kronecker delta function δ_{ij} such that the audio features for i^{th} timestamp attend to expression features at the j^{th} timestamp if and only if $i = j$. Mathematically,

$$\mathcal{M} = \delta_{ij} = \begin{cases} True & \text{if } i = j \\ False & \text{if } i \neq j \end{cases} \quad (7)$$

We demonstrate in the results (Section 6) that this alignment is crucial for learning audio-consistent expression codes. As a by-product, this further enables the model to generalize to arbitrary long audio signals during inference which we also show in our results. Our model is trained with the diffusion loss in Eq. 4, with input $\mathbf{x} = \{\theta_{exp}\}^{1:N}$ and conditioning $\mathbf{c} = \mathbf{A}^{1:N}$.

Expression Augmentation. Unlike other domains (like text-to-motion) where the synthesized motion can vary dramatically, speech-driven 3D facial animation is more constrained, requiring precise mouth alignment with the audio signal, making it prone to overfitting. To alleviate this, we propose a novel augmentation strategy capable of producing diverse expressions for a given audio signal. Our key insight is to augment the dataset to generate different expression codes for the same speech signal by randomly amplifying and suppressing its magnitude. Specifically, we randomly sample modulation factor r within the bounds $[a, b]$ as:

$$r \sim \text{Uniform}(a, b), \quad (8)$$

and scale the expression codes as $\{r \times \theta_{exp}\}^{1:N}$. We show in the results this augmentation strategy helps the model synthesize diverse expressions.

Sampling. FaceTalk uses a diffusion-based framework to learn to synthesize NPHM expression sequences of N frames $\{\theta_{exp}\}^{1:N} \in \mathbb{R}^{N \times 200}$. At each of the denoising timestep t , FaceTalk predicts the denoised sample and noises it back to timestamp $t - 1$:

$$\{\theta_{exp}\}_{t-1}^{1:N} = \mathcal{G}_\theta(\{\theta_{exp}\}_t^{1:N}, \mathbf{A}^{1:N}, t - 1), \quad (9)$$

terminating when it reaches $t = 0$. We train our model using classifier-free guidance [28]. We implement classifier-free guidance by randomly replacing the audio with null conditioning $\mathbf{A}^{1:N} = \Phi$ during training with 25% probability. During inference, we can use the weighted combination of conditional and unconditionally generated samples:

$$\{\theta_{exp}^c\}_t^{1:N} = w \cdot \{\theta_{exp}^c\}_t^{1:N} + (1 - w) \cdot \{\theta_{exp}^u\}_t^{1:N}, \quad (10)$$

where θ_{exp}^c and θ_{exp}^u refers to conditionally and unconditionally generated samples. To amplify the audio conditioning, we can use guidance strength $w > 1$.

Sequence Generation. The expression codes predicted above $\{\hat{\theta}_{exp}\}_t^{1:N}$ are then passed to the pretrained NPHM expression mapper $\{\mathcal{E}\}$ to obtain expression deformations $\{\delta_{exp}\}_t^{1:N}$. We then apply smoothing based on a Gaussian kernel to these deformations to remove the unwanted wiggle in the head and neck regions and to ensure that generating sequences are free from flickering artifacts. Specifically, we first define a control center \mathcal{C} on the mouth region in the canonical space as:

$$\mathcal{C} = [c_x, c_y, c_z]. \quad (11)$$

Next, using a 3D gaussian kernel with standard deviation $\Sigma = [\sigma_x, \sigma_y, \sigma_z]$, centered at \mathcal{C} , for each point sampled from 3D grid $\mathcal{P} \in \mathbb{R}^{|X| \times |Y| \times |Z|}$ as $\mathcal{P}_{xyz} = [p_x, p_y, p_z]$, we compute its distance from \mathcal{C} as:

$$d_{xyz} = \sqrt{\frac{(p_x - c_x)^2}{\sigma_x^2} + \frac{(p_y - c_y)^2}{\sigma_y^2} + \frac{(p_z - c_z)^2}{\sigma_z^2}}. \quad (12)$$

Based on this distance, smoothing weights are obtained from the Gaussian kernel with min-max normalization as:

$$w_{xyz} = \frac{1}{2\pi\sigma_x\sigma_y\sigma_z} \cdot e^{-\left(\frac{1}{2} \times d_{xyz}^2\right)} \quad (13)$$

$$w_{xyz}^k = \frac{w_{xyz}^k - \min(w_{xyz}^{1:M})}{\max(w_{xyz}^{1:M}) - \min(w_{xyz}^{1:M})}, \quad (14)$$

where $M = |X| \times |Y| \times |Z|$. The obtained expression deformations $\{\delta_{exp}\}_t^{1:N}$ are then multiplied with these smoothing weights $\{\mathbf{w}\}_t^{1:N} = \{(w^1, w^2, \dots, w^M)\}_t^{1:N}$ to obtain smoothed expression deformations as:

$$\{\delta_{exp}^s\}_t^{1:N} = \{\delta_{exp}\}_t^{1:N} \times \{\mathbf{w}\}_t^{1:N}. \quad (15)$$

These smoothed expression deformations $\{\delta_{exp}^s\}_t^{1:N}$, along with identity θ_{id} and predicted expression codes $\{\hat{\theta}_{exp}\}_t^{1:N}$ are passed to the NPHM identity MLP $\{I\}$ to obtain smoothed SDF, from which meshes are extracted using MC [42].

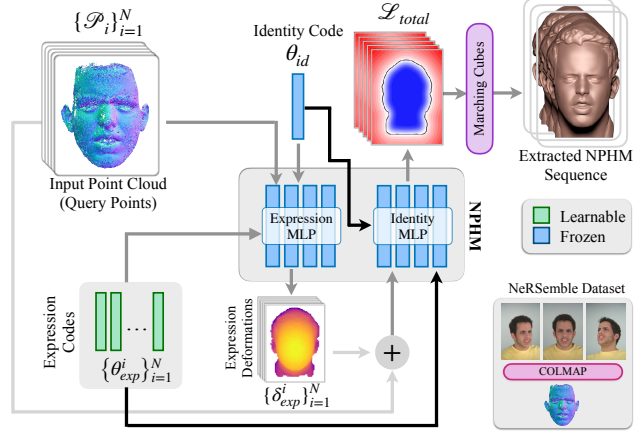


Figure 3. Given the pointcloud sequence $\{\mathcal{P}_i\}_{i=1}^N$ extracted from multi-view sequences from NeRSemble dataset [37] (bottom right), which also act as query points, we leverage the pretrained Expression MLP $\{\mathcal{E}\}$ to extract the expression deformations $\{\delta_{exp}^i\}_{i=1}^N$ and add them back to the input points to get the deformed points $\{\mathcal{P}'_i\}_{i=1}^N$. These points are then fed to the Identity MLP $\{\mathcal{I}\}$ which outputs the SDF. The expression codes $\{\theta_{exp}^i\}_{i=1}^N$ are optimized using overall loss \mathcal{L}_{total} . Note that both fixed identity code θ_{id} and learnable expression codes $\{\theta_{exp}^i\}_{i=1}^N$ are fed to both $\{\mathcal{I}\}$ and $\{\mathcal{E}\}$. Once optimized, the meshes are then extracted with Marching Cubes [42].

5. Dataset Creation

To train our latent diffusion model, we require temporally consistent NPHM expression codes as well as their corresponding audio. To this end, we leverage multi-view recordings of NeRSemble [37] captured with 16 cameras of talking human faces paired with corresponding audio signals. Notably, these sequences are captured for faces already present in the identity space of the NPHM model. Thus, for a given subject, only the expression codes corresponding to the multi-view sequences have to be estimated. Specifically, we first extract 3D pointclouds $\{\mathcal{P}_i\}_{i=1}^N$ from individual frames using COLMAP [54], where N refers to the number of frames in a sequence, with each $\mathcal{P}_i \in \mathbb{R}^{K \times 3}$ consisting of K points. Given the pretrained NPHM model, for the known facial shape θ_{id} , we then optimize the expression codes $\{\theta_{exp}^i\}_{i=1}^N$ to match the pointclouds. In particular, at every iteration, we randomly sample a subset \mathcal{S}_i of 5000 points from \mathcal{P}_i as $\mathcal{S}_i \sim \text{Uniform}(5000, \mathcal{P}_i)$, and then feed the sampled points \mathcal{S}_i , identity code θ_{id} and learnable expression code θ_{exp}^i into the expression decoder $\{\mathcal{E}\}$ to obtain the expression deformation δ_i for the sampled points \mathcal{S}_i . The deformation δ_i is then added to \mathcal{S}_i to obtain deformed points $\mathcal{D}_i = \mathcal{S}_i + \delta_i$. These deformed points \mathcal{D}_i are finally fed to the identity decoder $\{\mathcal{I}\}$ along with latent codes θ_{id} and θ_{exp}^i to obtain the SDF. The expression




Figure 4. Qualitative comparison for audio-driven face animation. Our approach maintains high fidelity while demonstrating rich mouth and nasolabial movements. In particular, we demonstrate more accurate lip articulation, precisely synchronized to phonetic movements.

codes are then optimized using the loss \mathcal{L}_{total} which is illustrated in Fig. 3. Naively optimizing the expression codes on a per-frame basis shows flickering artifacts. To mitigate this, we optimize expressions in groups of $n = 10$ frames in a sliding window fashion (with an overlap of 2 frames between adjacent windows) using an additional temporal regularization \mathcal{L}_{temp} . To further prevent the expression codes $\{\theta_{exp}^i\}_{i=1}^N$ from deviating too much from the distribution already learned by the NPHM model, we employ additional L2 expression regularization \mathcal{L}_{exp} . We minimize the objective \mathcal{L}_{total} defined as:

$$\mathcal{L}_{total} = \lambda_{sdf} \mathcal{L}_{sdf} + \lambda_{temp} \mathcal{L}_{temp} + \lambda_{reg} \mathcal{L}_{reg}. \quad (16)$$

The SDF loss \mathcal{L}_{sdf} , temporal regularization \mathcal{L}_{temp} , and expression regularization \mathcal{L}_{exp} can be defined as:

$$\mathcal{L}_{sdf} = \sum_{i=1}^{|n|} \frac{1}{|S|} \sum_{k=1}^{|S|} \|\text{SDF}(\mathcal{S}_i^k)\|_1 \quad (17)$$

$$\mathcal{L}_{temp} = \sum_{i=1}^{|n|} \|\theta_{exp}^{i+1} - \theta_{exp}^i\|_{\epsilon}, \quad \mathcal{L}_{exp} = \sum_{i=1}^{|n|} \|\theta_{exp}^i\|_2, \quad (18)$$

where $\|\cdot\|_1$, $\|\cdot\|_2$ and $\|\cdot\|_{\epsilon}$ refers to L1, L2 and Huber loss respectively, and n refers to number of frames simultaneously optimized. More details regarding aligning the coordinate system of NeRSemble dataset [37] with NPHM

can be found in the supplemental document. Once optimized, next to generate mesh sequences, we uniformly sample points from a 3D grid $\mathcal{P} \in \mathbb{R}^{|X| \times |Y| \times |Z|}$ such that $\mathcal{P}_{xyz} = [p_x, p_y, p_z]$ is a point in 3d space. These points, along with the identity θ_{id} and expression codes $\{\theta_{exp}^i\}_{i=1}^N$ are passed to the pretrained NPHM model outputting an SDF; meshes are then extracted with [42]. In total, our FaceTalk dataset consists of 1000 sequences, an order of magnitude larger than the existing datasets [14].

6. Results

We evaluate FaceTalk on the task audio-driven motion synthesis and compare it against state-of-the-art methods.

Evaluation Metrics. Unlike template-based methods that rely on LVE (Lip Vertex Error), our method produces SDFs for each frame; extracted mesh topologies thus vary across frames, making LVE inapplicable for evaluation. Thus, we employ LSE-D (Lip Sync Error Distance) [48] for quantitative evaluation of lip synchronization. For diversity, we compute FID and KID, as well as diversity scores [50]. As FID/KID may not reflect all quality considerations, when samples are limited, we include two established quality metrics: (1) FIQA (Face Image Quality Assessment) [43] (2) VQA (Video Quality Assessment) [74]. FIQA quantifies image quality for face recognition and similarity to in-the-wild real faces and VQA measures overall video qual-

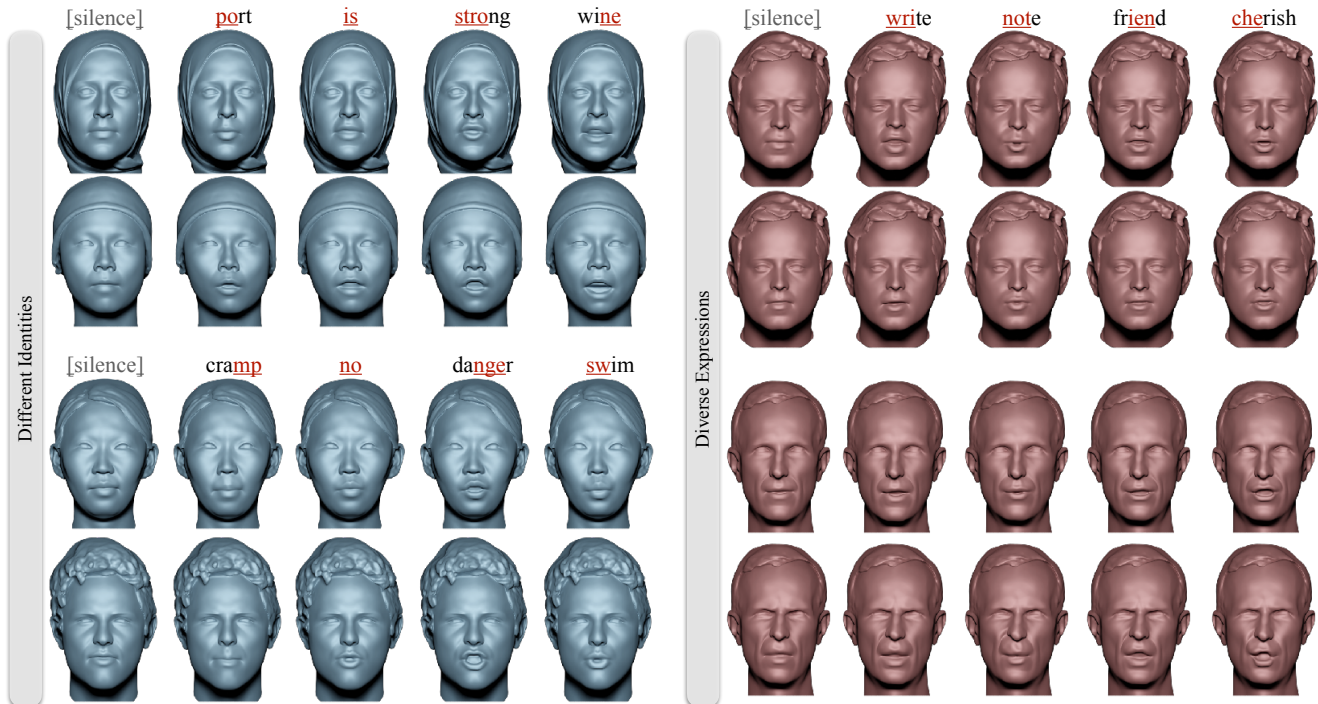


Figure 5. **Left.** Expressions generated by our method can easily be applied to diverse identities with complex geometry. **Right.** Given a speech signal, our method can further generate diversity in expression for the same identity. Note the difference in speaking style (intensity of mouth opening) as well as eyeblinks/frowning in the upper face area.

ity in terms of distortions/semantics/aesthetics. Finally, for perceptual evaluation, we perform a user study with 40 participants on a diverse set of 15 unseen audio clips. Most baseline methods are trained using Vocaset [14]. Thus, for a fair comparison, we construct a hybrid audio test set of 100 sequences, with 25 test audio clips each from Vocaset (~ 2 -3 seconds) and ours (~ 2 -3 seconds), and 50 (~ 5 -7 seconds) from LJSpeech [31]. Test audios were selected from identities not seen during training of any method. The metrics are reported on this hybrid test set.

Implementation Details. For dataset creation, we optimize NPHM expressions for 500 iterations in a group of $n = 10$ frames with a step size of 0.001 for iterations ≤ 300 , and 0.0001 otherwise. We use $\lambda_{sdf} = 10$, $\lambda_{temp} = 0.1$ and $\lambda_{reg} = 0.0025$ for Eq. 16 during optimization. To train our diffusion model, we randomly clip sequences to 2 seconds (48 frames) for efficient minibatching. We train with Adam optimizer with a learning rate of 0.0001. For diffusion, we use 1000 noising/denoising timestamps, a cosine schedule to add noise to input sequences. We apply data augmentation to the expression codes by uniformly modulating them.

Baseline Comparisons. This is the first work to perform speech-conditioned synthesis for volumetric head motion sequences. We compare with state-of-the-art template-based methods in Fig 4 and Tab. 1. Specifically, we compare against speech-driven animation methods [14, 21, 51, 75],

Method	LSE-D ↓	FID ↓	KID ↓	FIQA ↑	VQA ↑
VOCA [14]	13.6191	239.043	0.280	33.36	0.4251
MeshTalk [51]	12.9607	220.172	0.254	36.14	0.4855
FaceFormer [21]	11.9848	215.274	0.222	38.24	0.5227
CodeTalker [75]	11.8054	208.064	0.207	38.38	0.5274
Imitator [65]	11.6119	208.479	0.212	38.44	0.5384
EmoTalk [44]	11.7485	201.311	0.201	31.80	0.5037
EMOTE [15]	11.7192	227.924	0.247	39.44	0.5149
Ours	11.2737	40.692	0.009	45.75	0.6145

Table 1. In comparison to state of the art, FaceTalk more accurately matches the audio, while producing high perceptual fidelity.

personalized methods [65], as well as recent emotion-based methods [15, 44]. Our approach more accurately captures subtle mouth movements compared to less expressive baselines, resulting in better FID/KID. These scores were evaluated only on mouth region crops to avoid bias towards rest of facial geometry. Our method consistently achieves better lip-audio synchronization while also representing fine-scale facial details like creasing in the nasolabial folds with wider mouth motions. This is confirmed by our perceptual user study in Fig 6.

Generative Synthesis. Our generated expression codes are identity-agnostic, and are easily transferred to different identities, as shown in Fig. 5. Additionally, we can synthesize diverse expressions per-identity for a given audio.




Figure 6. User study comparison with baselines. We measure preference for (1) Overall Animation Quality, (2) Lip Synchronization and (3) Facial Realism. FaceTalk results are overwhelmingly preferred over the best baseline methods on all these aspects.

Our Method	LSE-D ↓	Diversity ↑
w/o expr. aug.	11.6229	1.61e-8
w/o facial smoothing	11.3488	0.34
w/o FiLM layer	12.029	0.006
w/o expr.-audio align.	13.720	5.17e-8
w/o diffusion	11.3217	0.000
Full (Ours)	11.2737	0.34

Table 2. Ablation over model design. Expression augmentation improves diversity, while facial smoothing alleviates inter-frame jitter. Using FiLM conditioning achieves accurate mouth pose and increased diversity. Without expression-audio alignment, the model ignores the audio signal. Without diffusion, the model fails to generate diverse results. Our full model with all components achieves the best results.

Architecture Ablations. We ablate our model design choices in Tab 2, with visual results shown in Fig 7, and refer to the supplemental for video demonstration. Our proposed expression augmentation helps the model to synthesize diverse expressions. Facial smoothing removes unwanted wobbliness in the head and neck region, thereby reducing the inter-frame jitter. FiLM conditioning ensures that expressions generated by our method are more pronounced and that expressions accurately match the audio signal. The expression-audio alignment mask is necessary to match the synthesized expression codes with the audio signal. Our diffusion training ensures that expressions generated by our method are high-quality and diverse.

Limitations. While FaceTalk effectively synthesizes high-fidelity and audio-synchronized facial expressions, it still has limitations. For instance, the use of a diffusion model requires multiple denoising steps during inference, limiting its real-time application. We believe that this could be improved by investigating efficient sampling techniques [72]. Currently, our method specializes in synthesizing only the expression codes. For holistic 3D facial animation, we need to extend its capability to synthesize facial identities. In the future, we would like to generate diverse identities aligned



Figure 7. Without FiLM conditioning, the model fails to synthesize accurate lip synchronization and even generates uncanny expressions. Without expression-audio alignment, the model fully ignores the audio signal synthesizing constant expression. Our full diffusion model synthesizes accurate lip articulation while maintaining temporal coherence which can be seen in the suppl. video.

with the gender inferred directly from the audio.

7. Conclusion

In this work, we have introduced FaceTalk, a new approach to synthesize animations of realistic volumetric human heads from audio. We introduced the first latent diffusion model for audio-driven head animation, producing significantly higher fidelity results than existing methods. FaceTalk leverages a parametric head model producing volumetric head representations to generate expressive, fine-grained motions such as eye blinks, skin creasing, etc. We also demonstrate the applicability of our method for other conditioning signals such as facial landmarks. We believe this is an important first step towards enabling the animation of highly detailed 3D face models, which can enable many new possibilities for content creation and digital avatars.

8. Acknowledgments

This work was supported by the ERC Starting Grant Scan2CAD (804724), the Bavarian State Ministry of Science and the Arts and coordinated by the Bavarian Research Institute for Digital Transformation (bidt), the German Research Foundation (DFG) Grant “Making Machine Learning on Static and Dynamic 3D Data Practical,” the German Research Foundation (DFG) Research Unit “Learning and Simulation in Visual Computing,” and Sony Semiconductor Solutions Corporation. We would like to thank Simon Giebenhain and Tobias Kirschstein for their help with dataset.

References

- [1] Chaitanya Ahuja and Louis-Philippe Morency. Language2pose: Natural language grounded pose forecasting, 2019. [1](#)
- [2] Simon Alexanderson, Rajmund Nagy, Jonas Beskow, and Gustav Eje Henter. Listen, denoise, action! audio-driven motion synthesis with diffusion models. *ACM Trans. Graph.*, 42(4):44:1–44:20, 2023. [1](#), [2](#)
- [3] Nikos Athanasiou, Mathis Petrovich, Michael J. Black, and Gül Varol. TEACH: Temporal Action Compositions for 3D Humans. In *International Conference on 3D Vision (3DV)*, 2022. [1](#)
- [4] Nikos Athanasiou, Mathis Petrovich, Michael J. Black, and Gül Varol. SINC: Spatial composition of 3D human motions for simultaneous action generation. In *International Conference on Computer Vision (ICCV)*, 2023. [1](#)
- [5] Samaneh Azadi, Akbar Shah, Thomas Hayes, Devi Parikh, and Sonal Gupta. Make-an-animation: Large-scale text-conditional 3d human motion generation, 2023. [2](#)
- [6] Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations, 2020. [2](#), [3](#), [4](#)
- [7] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 157–164, 2023. [1](#)
- [8] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [2](#)
- [9] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *arXiv*, 2021. [2](#)
- [10] Lele Chen, Zhiheng Li, Ross K. Maddox, Zhiyao Duan, and Chenliang Xu. Lip movements generation at a glance, 2018. [2](#)
- [11] 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.
- [12] Lele Chen, Guofeng Cui, Celong Liu, Zhong Li, Ziyi Kou, Yi Xu, and Chenliang Xu. Talking-head generation with rhythmic head motion. *arXiv preprint arXiv:2007.08547*, 2020.
- [13] Joon Son Chung, Amir Jamaludin, and Andrew Zisserman. You said that?, 2017. [2](#)
- [14] Daniel Cudeiro, Timo Bolkart, Cassidy Laidlaw, Anurag Ranjan, and Michael Black. Capture, learning, and synthesis of 3D speaking styles. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10101–10111, 2019. [2](#), [6](#), [7](#)
- [15] Radek Daněček, Kiran Chhatre, Shashank Tripathi, Yandong Wen, Michael Black, and Timo Bolkart. Emotional speech-driven animation with content-emotion disentanglement. In *SIGGRAPH Asia 2023 Conference Papers*, New York, NY, USA, 2023. Association for Computing Machinery. [2](#), [7](#)
- [16] Dipanjan Das, Sandika Biswas, Sanjana Sinha, and Brojeshwar Bhowmick. Speech-driven facial animation using cascaded gans for learning of motion and texture. In *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX*, page 408–424, Berlin, Heidelberg, 2020. Springer-Verlag. [2](#)
- [17] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis, 2021. [2](#)
- [18] Rohan Dhesikan and Vignesh Rajmohan. Sketching the future (stf): Applying conditional control techniques to text-to-video models, 2023. [2](#)
- [19] Michail Christos Doukas, Stefanos Zafeiriou, and Viktoriia Sharmanska. Headgan: One-shot neural head synthesis and editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14398–14407, 2021. [2](#)
- [20] Fanda Fan, Chaoxu Guo, Litong Gong, Biao Wang, Tiezheng Ge, Yuning Jiang, Chunjie Luo, and Jianfeng Zhan. Hierarchical masked 3d diffusion model for video outpainting, 2023. [2](#)
- [21] Yingruo Fan, Zhaojiang Lin, Jun Saito, Wenping Wang, and Taku Komura. Faceformer: Speech-driven 3d facial animation with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [2](#), [7](#)
- [22] Guy Gafni, Justus Thies, Michael Zollhöfer, and Matthias Nießner. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8649–8658, 2021. [2](#)
- [23] Simon Giebenhain, Tobias Kirschstein, Markos Georgopoulos, Martin Rünz, Lourdes Agapito, and Matthias Nießner. Learning neural parametric head models. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2023. [2](#), [3](#), [4](#)
- [24] Simon Giebenhain, Tobias Kirschstein, Markos Georgopoulos, Martin Rünz, Lourdes Agapito, and Matthias Nießner. MonoNPHM: Dynamic head reconstruction from monocular videos. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2024. [2](#), [3](#)
- [25] Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and Li Cheng. Action2motion: Conditioned generation of 3d human motions. In *Proceedings of the 28th ACM International Conference on Multimedia*, page 2021–2029, New York, NY, USA, 2020. Association for Computing Machinery. [1](#)
- [26] Yudong Guo, Keyu Chen, Sen Liang, Yongjin Liu, Hujun Bao, and Juyong Zhang. Ad-nerf: Audio driven neural radiance fields for talking head synthesis. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. [2](#)
- [27] Siddharth Gururani, Arun Mallya, Ting-Chun Wang, Rafael Valle, and Ming-Yu Liu. Space: Speech-driven portrait animation with controllable expression, 2022. [2](#)
- [28] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. [5](#)

- [29] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *arXiv preprint arxiv:2006.11239*, 2020. 2
- [30] Rongjie Huang, Max W. Y. Lam, Jun Wang, Dan Su, Dong Yu, Yi Ren, and Zhou Zhao. Fastdiff: A fast conditional diffusion model for high-quality speech synthesis, 2022. 2
- [31] Keith Ito and Linda Johnson. The lj speech dataset. <https://keithito.com/LJ-Speech-Dataset/>, 2017. 7
- [32] 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 (CVPR)*, 2021. 2
- [33] Xinya Ji, Hang Zhou, Kaisiyuan Wang, Qianyi Wu, Wayne Wu, Feng Xu, and Xun Cao. Eamm: One-shot emotional talking face via audio-based emotion-aware motion model. In *ACM SIGGRAPH 2022 Conference Proceedings*, 2022. 2
- [34] Prajwal K R, Rudrabha Mukhopadhyay, Jerin Philip, Abhishek Jha, Vinay Namboodiri, and C V Jawahar. Towards automatic face-to-face translation. In *Proceedings of the 27th ACM International Conference on Multimedia*, page 1428–1436, New York, NY, USA, 2019. Association for Computing Machinery. 2
- [35] 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 Trans. Graph.*, 36(4), 2017. 2
- [36] Jihoon Kim, Jiseob Kim, and Sungjoon Choi. Flame: Free-form language-based motion synthesis & editing. *arXiv preprint arXiv:2209.00349*, 2022. 1, 2
- [37] Tobias Kirschstein, Shenhan Qian, Simon Giebenhain, Tim Walter, and Matthias Nießner. Nersemble: Multi-view radiance field reconstruction of human heads. *ACM Trans. Graph.*, 42(4), 2023. 2, 5, 6
- [38] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis, 2021. 2
- [39] Max WY Lam, Jun Wang, Dan Su, and Dong Yu. Bddm: Bilateral denoising diffusion models for fast and high-quality speech synthesis. In *International Conference on Learning Representations*, 2022. 2
- [40] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 36(6):194:1–194:17, 2017. 1
- [41] Xian Liu, Yinghao Xu, Qianyi Wu, Hang Zhou, Wayne Wu, and Bolei Zhou. Semantic-aware implicit neural audio-driven video portrait generation. *arXiv preprint arXiv:2201.07786*, 2022. 2
- [42] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, page 163–169, New York, NY, USA, 1987. Association for Computing Machinery. 3, 5, 6
- [43] Fu-Zhao Ou, Xingyu Chen, Ruixin Zhang, Yuge Huang, Shaoxin Li, Jilin Li, Yong Li, Liujuan Cao, and Yuan-Gen Wang. SDD-FIQA: Unsupervised face image quality assessment with similarity distribution distance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 6
- [44] Ziqiao Peng, Haoyu Wu, Zhenbo Song, Hao Xu, Xiangyu Zhu, Jun He, Hongyan Liu, and Zhaoxin Fan. Emotalk: Speech-driven emotional disentanglement for 3d face animation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20687–20697, 2023. 2, 7
- [45] Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer, 2017. 3, 4
- [46] Mathis Petrovich, Michael J. Black, and Gül Varol. Action-conditioned 3D human motion synthesis with transformer VAE. In *International Conference on Computer Vision (ICCV)*, 2021. 1
- [47] Mathis Petrovich, Michael J. Black, and Gül Varol. TEMOS: Generating diverse human motions from textual descriptions. In *European Conference on Computer Vision (ECCV)*, 2022. 1
- [48] K R Prajwal, Rudrabha Mukhopadhyay, Vinay P. Namboodiri, and C.V. 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*, page 484–492, New York, NY, USA, 2020. Association for Computing Machinery. 2, 6
- [49] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents, 2022. 2
- [50] Zhiyuan Ren, Zhihong Pan, Xin Zhou, and Le Kang. Diffusion motion: Generate text-guided 3d human motion by diffusion model, 2023. 6
- [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/CVF International Conference on Computer Vision (ICCV)*, pages 1173–1182, 2021. 2, 7
- [52] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [53] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding, 2022. 2
- [54] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5
- [55] Shuai Shen, Wanhua Li, Zheng Zhu, Yueqi Duan, Jie Zhou, and Jiwen Lu. Learning dynamic facial radiance fields for few-shot talking head synthesis. In *ECCV*, 2022. 2
- [56] Shuai Shen, Wenliang Zhao, Zibin Meng, Wanhua Li, Zheng Zhu, Jie Zhou, and Jiwen Lu. Difftalk: Crafting diffusion

- models for generalized audio-driven portraits animation. In *CVPR*, 2023. 2
- [57] Linsen Song, Wayne Wu, Chen Qian, Ran He, and Chen Change Loy. Everybody’s talkin’: Let me talk as you want, 2020. 2
- [58] Stefan Stan, Kazi Injamamul Haque, and Zerrin Yumak. Facediffuser: Speech-driven 3d facial animation synthesis using diffusion. In *Proceedings of the 16th ACM SIGGRAPH Conference on Motion, Interaction and Games*, New York, NY, USA, 2023. Association for Computing Machinery. 2
- [59] Michał Stypułkowski, Konstantinos Vougioukas, Sen He, Maciej Zikeba, Stavros Petridis, and Maja Pantic. Diffused Heads: Diffusion Models Beat GANs on Talking-Face Generation. In <https://arxiv.org/abs/2301.03396>, 2023. 2
- [60] Zhiyao Sun, Tian Lv, Sheng Ye, Matthieu Gaetan Lin, Jenny Sheng, Yu-Hui Wen, Minjing Yu, and Yong jin Liu. Diffposetalk: Speech-driven stylistic 3d facial animation and head pose generation via diffusion models, 2023. 2
- [61] Supasorn Suwajanakorn, Steven M. Seitz, and Ira Kemelmacher-Shlizerman. Synthesizing obama: Learning lip sync from audio. *ACM Trans. Graph.*, 2017. 2
- [62] Anni Tang, Tianyu He, Xu Tan, Jun Ling, Runnan Li, Sheng Zhao, Li Song, and Jiang Bian. Memories are one-to-many mapping alleviators in talking face generation. *arXiv preprint arXiv:2212.05005*, 2022. 2
- [63] Jiapeng Tang, Angela Dai, Yinyu Nie, Lev Markhasin, Justus Thies, and Matthias Niessner. Dphms: Diffusion parametric head models for depth-based tracking. 2024. 2
- [64] Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. Human motion diffusion model. In *The Eleventh International Conference on Learning Representations*, 2023. 1, 2
- [65] Balamurugan Thambiraja, Ikhsanul Habibie, Sadegh Aliakbarian, Darren Cosker, Christian Theobalt, and Justus Thies. Imitator: Personalized speech-driven 3d facial animation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20621–20631, 2023. 2, 7
- [66] Justus Thies, Mohamed Elgharib, Ayush Tewari, Christian Theobalt, and Matthias Nießner. Neural voice puppetry: Audio-driven facial reenactment. *ECCV 2020*, 2020. 2
- [67] Jonathan Tseng, Rodrigo Castellon, and C Karen Liu. Edge: Editable dance generation from music. *arXiv preprint arXiv:2211.10658*, 2022. 1, 2
- [68] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2023. 3, 4
- [69] Vikram Voleti, Alexia Jolicoeur-Martineau, and Christopher Pal. Mcvd: Masked conditional video diffusion for prediction, generation, and interpolation. In *(NeurIPS) Advances in Neural Information Processing Systems*, 2022. 2
- [70] Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. End-to-end speech-driven facial animation with temporal gans, 2018. 2
- [71] Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Realistic speech-driven facial animation with gans, 2019. 2
- [72] Daniel Watson, Jonathan Ho, Mohammad Norouzi, and William Chan. Learning to efficiently sample from diffusion probabilistic models, 2021. 8
- [73] O. Wiles, A.S. Koepke, and A. Zisserman. X2face: A network for controlling face generation by using images, audio, and pose codes. In *European Conference on Computer Vision*, 2018. 2
- [74] Haoning Wu, Erli Zhang, Liang Liao, Chaofeng Chen, Jingwen Hou Hou, Annan Wang, Wenxiu Sun Sun, Qiong Yan, and Weisi Lin. Exploring video quality assessment on user generated contents from aesthetic and technical perspectives. In *International Conference on Computer Vision (ICCV)*, 2023. 6
- [75] Jinbo Xing, Menghan Xia, Yuechen Zhang, Xiaodong Cun, Jue Wang, and Tien-Tsin Wong. Codetalker: Speech-driven 3d facial animation with discrete motion prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12780–12790, 2023. 2, 7
- [76] Shunyu Yao, RuiZhe Zhong, Yichao Yan, Guangtao Zhai, and Xiaokang Yang. Dfa-nerf: Personalized talking head generation via disentangled face attributes neural rendering. *arXiv preprint arXiv:2201.00791*, 2022. 2
- [77] Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Shaoli Huang, Yong Zhang, Hongwei Zhao, Hongtao Lu, and Xi Shen. T2m-gpt: Generating human motion from textual descriptions with discrete representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1
- [78] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023. 2
- [79] Hang Zhou, Yu Liu, Ziwei Liu, Ping Luo, and Xiaogang Wang. Talking face generation by adversarially disentangled audio-visual representation. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2019. 2
- [80] Yang Zhou, Xintong Han, Eli Shechtman, Jose Echevarria, Evangelos Kalogerakis, and Dingzeyu Li. Makeltalk: Speaker-aware talking-head animation. *ACM Trans. Graph.*, 39(6), 2020. 2
- [81] Lingting Zhu, Xian Liu, Xuanyu Liu, Rui Qian, Ziwei Liu, and Lequan Yu. Taming diffusion models for audio-driven co-speech gesture generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10544–10553, 2023. 2