

# FLOAT: Generative Motion Latent Flow Matching for Audio-driven Talking Portrait

Taekyung Ki<sup>1\*</sup> Dongchan Min<sup>1</sup> Gyeongsu Chae<sup>2</sup>

<sup>1</sup>KAIST <sup>2</sup>DeepBrain AI Inc.

{taekyung.ki, alsehdcks95}@kaist.ac.kr gc@deepbrain.io

<https://deepbrainai-research.github.io/float/>




Figure 1. FLOAT can generate a talking portrait video from a single source image and audio where the talking motion is generated by the motion latent flow matching. It can enhance the emotion-related talking motion by leveraging speech-driven emotion labels, a natural way of emotion-aware motion control.

## Abstract

With the rapid advancement of diffusion-based generative models, portrait image animation has achieved remarkable results. However, it still faces challenges in temporally consistent video generation and fast sampling due to its iterative sampling nature. This paper presents FLOAT, an audio-driven talking portrait video generation method based on flow matching generative model. Instead of a pixel-based latent space, we take advantage of a learned orthogonal motion latent space, enabling efficient generation and editing of temporally consistent motion. To achieve this, we introduce a transformer-based vector field predictor with an effective frame-wise conditioning mechanism. Additionally, our method supports speech-driven emotion enhancement, enabling a natural incorporation of expressive motions. Extensive experiments demonstrate that our method outperforms state-of-the-art audio-driven talking portrait methods in terms of visual quality, motion fidelity, and efficiency.

\*This work was done during South Korea Mandatory Military Service at DeepBrain AI Inc.

## 1. Introduction

Animating a single image using a driving audio (*i.e.*, audio-driven talking portrait generation) has gained significant attention in recent years for its great potential in avatar creation, video conferencing, virtual avatar chat, and user-friendly customer service. It aims to synthesize natural talking motion from audio signals, including accurate lip synchronization, rhythmical head movements, and fine-grained facial expressions. However, generating such motion solely from audio is extremely challenging due to its one-to-many correlation between audio and motion. In the earlier stage of this field, many works [8, 19, 28, 47, 51, 86] focus on generating accurate lip movements by relying on learned audio-lip alignment losses [9, 45].

To comprehensively extend the range of motion, some works [45, 66, 84] incorporate probabilistic generative models, such as VAE [29] and normalizing flow [53], turning the motion generation into probabilistic sampling. However, these models still lack expressiveness in generated motion due to the limited capacity of these generative models.

Recent talking portrait generation methods [7, 21, 26, 36,

44, 63, 68, 71, 76, 79], powered by diffusion-based generative models [23, 61], successfully mitigate this expressiveness issue. EMO [68] introduces a promising approach to this field [7, 26, 71, 76, 79] by employing a strong pre-trained image diffusion model (*i.e.*, StableDiffusion [54]) and lifting it into video generation [24]. However, it still faces challenges in generating temporally coherent videos and achieving sampling efficiency, requiring tens of minutes for a few seconds of video. Moreover, they heavily rely on auxiliary facial prior, such as bounding boxes [68, 79], 2D landmarks and skeletons [7, 26, 82], or 3D meshes [76], which significantly restricts the diversity and the fidelity of head movements due to their strong spatial bias.

In this paper, we present FLOAT, an audio-driven talking portrait video generation model based on flow matching generative model in a motion latent space. Flow matching [35, 37] has emerged as a promising alternative to diffusion models due to its fast and high-quality sampling. By modeling talking motion within a learned motion latent space [75], we can more efficiently sample temporally consistent motion latents. This is achieved by a simple yet effective transformer-based [70] vector field predictor, inspired by DiT [48]. Since our motion latent space has orthogonal structure, our method can manipulate head motion of the generated video using its basis. Furthermore, our method supports natural emotion-aware motion enhancement driven by speech. Our contributions are summarized as follows:

- We present, **FLOAT**, flow matching based audio-driven talking portrait generation model using a learned orthogonal motion latent space, enabling to generate talking portrait videos with reduced sampling steps.
- We introduce a simple yet effective transformer-based flow vector field predictor for temporally consistent motion latent sampling, which also enables the speech-driven emotional controls.
- Extensive experiments demonstrate that FLOAT achieves state-of-the-art performance compared to both diffusion- and non-diffusion-based methods.

## 2. Related Works

### 2.1. Diffusion Models and Flow Matching

**Diffusion Models** Diffusion models or score-based generative models [13, 23, 46, 54, 60, 61] are generative models that gradually diffuse input signals into Gaussian noise and learn the denoising reverse process for the generative modeling. They have shown remarkable results in various generation tasks, such as unconditional image and video generation [4, 16, 48], text-to-image generation [52, 54, 55], text-to-video generation [4, 20], conditional image generation [24, 82], and 3D human generation [31, 64, 67].

**Accelerating Diffusion Models** While diffusion models

demonstrate superior performance, their iterative sampling nature still bottlenecks the efficient generation compared to VAEs [29], normalizing flow [53], and GANs [18]. To overcome this limitation, several works have been developed to boost the sampling speed of the diffusion models. StableDiffusion (SD) [54] partially mitigates this problem by moving the diffusion process from the pixel space to the spatial latent space, establishing itself as a pivotal framework among diffusion models. Another line of research has developed the sampling solvers [40, 41] based on ordinary differential equations (ODEs). Meanwhile, model distillation [22] has been introduced to transfer the knowledge of the learned diffusion models into a student model, enabling one (or a few) steps of generation [27, 34, 38, 42, 62]. However, these approaches involve substantial effort to create a well-trained diffusion model and suffer from training instability.

**Flow Matching** Flow matching [35, 37] stands out as an alternative to diffusion models for its high sampling speed and competitive sample quality compared to diffusion models [10, 17, 32, 35, 50]. It belongs to the family of flow-based generative models, which estimates a transformation (referred to as a *flow*) between a prior distribution (*e.g.*, Gaussian) and a target distribution. Unlike the normalizing flow [14, 53] that directly estimates the noise-to-data transformation under specific architectural constraints (*e.g.*, affine coupling), flow matching regresses the time-dependent vector field that *generates* this flow by solving its corresponding ODEs [6] with flexible architectures. One specific design of flow matching is an optimal transport (OT) based one, which transforms the data distribution along the straight path with constant velocity [35].

Our audio-driven talking portrait method employs flow matching to generate the natural talking motions. Thanks to the architectural flexibility of flow matching, we use transformer-encoder architecture [70] to estimate the generating vector field, allowing us to take the video temporal consistency into account.

### 2.2. Audio-driven Portrait Animation

Audio-driven portrait animation is the task of generating a realistic talking portrait video using a single portrait image and driving audio [45, 73, 84, 87, 88]. Since audio-to-motion relation is basically a one-to-many problem, several works utilize additional facial prior for driving conditions, *e.g.*, 2D facial landmarks [7, 21, 26, 71, 76, 88], 3D prior [8, 43, 44, 81, 84], or emotional labels [25, 65, 80]. In earlier stages, most works [8, 19, 28, 51] focused on generating accurate lip motion from audio by utilizing the lip-sync discriminator [9]. These approaches have advanced to generating audio-related head poses in a probabilistic way. For example, StyleTalker [45] uses normalizing flow [14, 53] to generate the head motion from audio, while SadTalker [84] uses audio-conditional variational inference [29] to learn

the 3DMM coefficients [2], bridging the intermediate representations of a pre-trained portrait animator [74].

Meanwhile, several works [25, 65, 72, 77] focus on an emotion-aware talking portrait generation. In particular, EAMM [25] considers an emotion as the complementary displacement of facial motion, and learns these displacement from an emotion label extracted from the image.

Recent audio-driven talking portrait methods powered by diffusion models show remarkable results [7, 26, 36, 44, 68, 71, 76, 79, 80]. Specifically, EMO [68] and subsequent extensions [7, 71, 76, 79] utilize the pre-trained SD [54] as their backbone to leverage generative prior trained on the large-scale image datasets. They introduce additional modules, *e.g.*, ReferenceNet [24] and Temporal Transformer [20], to preserve input identity and enhance the video temporal consistency, respectively. However, these modules introduces additional computational cost, requiring several minutes for a few seconds of video, and still suffer from video-level artifacts, such as noisy frames, and flickering.

VASA-1 [80] addresses the sampling time issue by sampling motion latents [15], producing lifelike talking portraits. Our method takes advantage of this approach. However, unlike [80], our motion latent space has a strong linear orthogonal structure represented by a computable basis, enabling to manipulate the generated motion at the test-time without external driving signals. Based on this orthogonality, we employ OT-based flow matching for motion latent sampling along a straight line with reduced sampling steps.

### 3. Preliminaries: (Conditional) Flow Matching

Let  $x \in \mathbb{R}^d$  be a data,  $t \in [0, 1]$  be the time, and  $q$  be an unknown target distribution. We can define a *flow* as a time-dependent transformation  $\varphi_t : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$  that transforms a tractable prior distribution  $p_0$  to the distribution  $p_1 \approx q$ . This flow  $\varphi_t$  further introduces a *probability flow path*  $p_t : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}_{>0}$  and a *generating vector field*  $v_t : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$  where  $p_t$  is defined by the push-forwarding

$$p_t(x) = p_0(\varphi_t^{-1}(x)) \det \left| \frac{\partial \varphi_t^{-1}(x)}{\partial x} \right|, \quad (1)$$

and  $v_t$  generates  $\varphi_t$  by means of an ordinary differential equation (ODE) [6]:

$$\frac{d}{dt} \varphi_t(x) = v_t(\varphi_t(x)) \quad \text{and} \quad \varphi_0(x) = x. \quad (2)$$

Flow matching [35] aims to estimate the target generating vector field  $u_t$  with a neural network parameterized by  $\theta$ :

$$\mathcal{L}_{\text{FM}}(\theta) := \|v_t(x; \theta) - u_t(x)\|_2^2, \quad (3)$$

where  $t \sim \mathcal{U}[0, 1]$  and  $x \sim p_t(x)$ . However, the target generating vector field  $u_t$  and the sample distribution  $p_t$  are

intractable. To address this issue, [35] proposes a method for constructing a “conditional” probability path  $p_t(\cdot|x_1)$  as well as target “conditional” vector field  $u_t(\cdot|x_1)$  using a sample  $x_1 \sim q$  as a condition. And they prove that the following objective

$$\mathcal{L}_{\text{CFM}}(\theta) := \|v_t(x; \theta) - u_t(x|x_1)\|_2^2, \quad (4)$$

where  $t \sim \mathcal{U}[0, 1]$  and  $x \sim p_t(x|x_1)$ , is equivalent to (3) with respect to the gradient  $\nabla_\theta$ .

One natural way of constructing  $u_t(\cdot|x_1)$  is a “straight line” that connects  $x_0 \sim p_0$  and  $x_1 \sim q$ , drawing an *optimal transport (OT)* path with constant velocity [35]. Specifically, a linear time interpolation between  $x_0$  and  $x_1$  gives us the flow  $x_t = \varphi_t(x) = (1-t)x_0 + tx_1$ , the conditional probability path  $p_t(x|x_1)$  defined via the affine transformation  $p_t(x|x_1) = \mathcal{N}(x|tx_1, (1-t)^2I)$ , and the target generating vector field  $u_t(x|x_1) = x_1 - x_0$ . This specific choice turns the objective (4) into

$$\mathcal{L}_{\text{OT}}(\theta) := \|v_t((1-t)x_0 + tx_1; \theta) - (x_1 - x_0)\|_2^2, \quad (5)$$

where  $t \sim \mathcal{U}[0, 1]$ ,  $x_0 \sim p_0$ , and  $x_1 \sim q$ , all of which are tractable.

**Classifier-free Vector Field** [10] formulates a classifier-free vector field (CFV) technique for flow matching, which enables class-conditional sampling more controllable manner without any extra classifier trained on noisy trajectory. Formally, CFV compute the modified vector field  $\tilde{v}_t$  by

$$\tilde{v}_t(x_t, c; \theta) \approx \gamma v_t(x_t, c; \theta) + (1-\gamma)v_t(x_t, c = \emptyset; \theta), \quad (6)$$

where  $\gamma$  denotes the guidance scale.  $v_t(x_t, c = \emptyset; \theta)$  is the predicted vector field without a driving condition  $c$ . For more details, please refer to [10, 35].

### 4. Method: Flow Matching for Audio-driven Talking Portrait

We provide an overview of FLOAT in Fig. 2. Given source image  $S \in \mathbb{R}^{3 \times H \times W}$ , and a driving audio signal  $a^{1:L} \in \mathbb{R}^{L \times d_a}$  of length  $L$ , our method generates a video

$$\hat{D}^{1:L} = (\hat{D}^l)_{l=1}^L \in \mathbb{R}^{L \times 3 \times H \times W} \quad (7)$$

of  $L$  frames, featuring audio-synchronized talking head motions, including both verbal and non-verbal motions. Our method consists of two phases. First, we pre-train a motion auto-encoder, which provides us with the expressive and smooth motion latent space for the talking portraits (Sec. 4.1). Next, we employ OT-based flow matching [35] to generate a sequence of motion latents with a transformer-based vector field predictor using the driving audio, which is decoded to the talking portrait videos (Sec. 4.2). We also incorporate speech-driven emotions as the driving conditions, achieving automatic emotion-aware talking portrait generation without any extra user input for emotion.




Figure 2. Overview of FLOAT. We encode the source image  $S \in \mathbb{R}^{3 \times H \times W}$  into the latent with the explicit identity-motion decomposition  $w_s = w_{s \rightarrow r} + w_{r \rightarrow s} \in \mathbb{R}^d$ . Given audio segments  $a^{-L':L} \in \mathbb{R}^{(L'+L) \times d_a}$  of the length  $L' + L$  and the reference motion  $w_{r \rightarrow s} \in \mathbb{R}^d$ , and the speech-driven emotion label  $w_e \in \mathbb{R}^7$ , a flow matching transformer estimates the generating vector field  $v_t(\varphi_t(x_0), \mathbf{c}_t; \theta) \in \mathbb{R}^{L \times d}$  from noisy motion latents, which is used to solve corresponding ODE and generates the motion latents  $w_{r \rightarrow \hat{D}^{1:L}}$ . Finally, the sequence of latents  $w_{S \rightarrow \hat{D}^{1:L}} := (w_{S \rightarrow r} + w_{r \rightarrow \hat{D}^t})_{t=1}^L$  are decoded into the video  $\hat{D}^{1:L} \in \mathbb{R}^{L \times 3 \times H \times W}$ .

#### 4.1. Motion Latent Auto-encoder

Recent talking portrait methods utilize the VAE of StableDiffusion (SD) [54] due to its rich semantic pixel-based latent space. However, they often struggle to generate temporally consistent frames when lifted to video generating tasks [7, 24, 68, 79, 89]. Thus, our first goal for realistic talking portrait is to obtain *good* motion latent space, capturing both global (e.g., head motion) and fine-grained local (e.g., facial expressions, mouth and pupil movement) dynamics.

Instead of VAE of SD, we pre-train a motion latent auto-encoder that encodes images into motion latents. This is done by training an auto-encoder to reconstruct a driving image from a source image sampled from the same video clip, enforcing the encoder to implicitly capture both temporally adjacent and distant motions. Inspired by LIA [75], we introduce a learned orthonormal basis that decomposes the motion along distinct orthogonal directions. Specifically, our motion auto-encoder encodes the source  $S$  into the latent  $w_S \in \mathbb{R}^d$  with following explicit decomposition:

$$w_S := w_{S \rightarrow r} + w_{r \rightarrow S}, \quad (8)$$

where  $w_{S \rightarrow r} \in \mathbb{R}^d$  is the identity latent and

$$w_{r \rightarrow S} = \sum_{m=1}^M \lambda_m(S) \cdot \mathbf{v}_m \in \mathbb{R}^d \quad (9)$$

is the motion latent with  $\lambda(S) := (\lambda_m(S))_{m=1}^M \in \mathbb{R}^M$  being the source-dependent motion coefficients that span the learned source-agnostic motion basis  $V := \{\mathbf{v}_m\}_{m=1}^M \subseteq \mathbb{R}^d$ . In this space,  $\lambda_m(S)$  is the intensity of the motion direction  $\mathbf{v}_m$ .

As shown in Fig. 6, our method enables motion editing of the sampled (generated) motion using only the basis  $V$  and its orthogonality, as stated in Eq. (15).

**Improving Fidelity of Facial Components:**  $\mathcal{L}_{\text{comp-lp}}$  The expressiveness of generated motions and the image fidelity




Figure 3. Efficacy of  $\mathcal{L}_{\text{comp-lp}}$  for fine-grained motion and fidelity.

are determined by the motion space and the motion auto-encoder. However, as resolution increases, fine details in small facial regions (e.g., teeth, eyeballs) often get buried in large-scale dynamics. To address this issue, we propose a *facial component perceptual loss*  $\mathcal{L}_{\text{comp-lp}}$  using [59, 83] that significantly improves the image fidelity (e.g., teeth and eyes) as well as fine-grained motions (e.g., eyeball and eyebrows movements). As shown in Fig. 3,  $\mathcal{L}_{\text{comp-lp}}$  allows us to generate high-fidelity facial components and their fine-grained motions without relying on pre-trained foundation models, such as StableDiffusion [54].

#### 4.2. Flow Matching in Motion Latent Space

Armed with this linear orthogonal space, we employ OT-based flow matching [35, 37] for the motion sampling. Specifically, we predict a vector field  $v_t(x_t, \mathbf{c}_t; \theta) \in \mathbb{R}^{L \times d}$  where  $x_t$  is the sample at flow time  $t \in [0, 1]$ , and  $\mathbf{c}_t \in \mathbb{R}^{L \times h}$  represents the driving conditions for  $L$  consequent frames. This vector field generates the flow  $\varphi_t : [0, 1] \times \mathbb{R}^{L \times d} \rightarrow \mathbb{R}^{L \times d}$  of  $L$  frames by solving ODE (Eq. (2)). As illustrated in Fig. 4, we build our vector field predictor upon the transformer encoder [70] architecture. Specifically, we adopt DiT [48] architecture, but decouple frame-wise conditioning from time-axis attention mechanism, which enables us to model temporally consistent motion latents.

In DiT [48], distinct semantic tokens are modulated by a




Figure 4. Frame-wise vector field predictor block at inference.

single diffusion time step embedding and class embedding through adaptive layer normalization (AdaLN). In contrast, our vector field predictor modulates each  $l$ -th input latent with its corresponding  $l$ -th condition and then combines their temporal relations through a masked self-attention layer that attends to  $2 \cdot T$  neighboring frames. Formally, for each  $l$ -th frame, frame-wise AdaLN and frame-wise gating are computed by

$$\gamma_i^l \times \text{LN}(X_t^l) + \beta_i^l \in \mathbb{R}^h \quad \text{and} \quad \alpha_i^l \times X_t^l \in \mathbb{R}^h, \quad (10)$$

respectively, where  $i \in \{1, 2\}$ ,  $h$  is the hidden dimension,  $\text{LN}(\cdot)$  denotes layer norm [33], and  $X_t^l$  is the  $l$ -th input for each operation at flow time  $t \in [0, 1]$ . The coefficients  $\alpha_i^l, \beta_i^l, \gamma_i^l \in \mathbb{R}^h$  are computed from the condition  $\mathbf{c}_t^l \in \mathbb{R}^h$  through a linear layer, *ToScaleShift*, as depicted in Fig. 4.

**Speech-driven Emotion Enhancement** *How can we make talking motions more expressive and natural?* During talking, humans naturally reflect their emotions through their voices, and these emotions influence talking motions. For instance, a person who speaks sadly may be more likely to shake the head and avoid eye contact. This non-verbal motion derived from emotions crucially impacts the naturalness of a talking portrait.

Existing works [25, 72, 80] use image-emotion paired data or image-driven emotion predictor [56] to generate the emotion-aware motion. In contrast, we incorporate speech-driven emotions, a more intuitive way of controlling emotion for audio-driven talking portrait. Specifically, we utilize a pre-trained speech emotion predictor [49] that produces softmax probabilities of seven distinct emotions: *angry, disgust, fear, happy, neutral, sad, and surprise*, which we then input into the vector field predictor.

However, as people do not always speak with a single, clear emotion, determining emotions solely from audio is often ambiguous [25]. Naive introduction of speech-driven emotion can make emotion-aware motion generation more challenging. To address this issue, we inject the emotions together with other driving conditions at training phase and modify them at inference phase.

**Driving Conditions** We concatenate the audio representation  $a^{1:L} \in \mathbb{R}^{L \times d_a}$  of a pre-trained Wav2Vec2.0 [1], the speech emotion label  $w_e \in \mathbb{R}^7$ , and the source motion latent  $w_{r \rightarrow S} \in \mathbb{R}^d$ . Next, we add the flow time step embedding

$\text{Emb}(t) \in \mathbb{R}^h$  to these conditions, producing  $\mathbf{c}_t \in \mathbb{R}^{L \times h}$  via a linear layer, *ToCondition*, as depicted in Fig. 2, where  $\text{Emb}(t)$  is computed using the sinusoidal position embedding [70].

**Training** We train FLOAT by reconstructing a target vector field computed from driving frames using the corresponding audio segments and a source motion latent. We choose a pair of driving motions and corresponding audio ( $w_{r \rightarrow D^{1:L}}, a^{1:L}$ ), and construct the target vector field  $u_t(x|w_{r \rightarrow D^{1:L}}) = w_{r \rightarrow D^{1:L}} - x_0 \in \mathbb{R}^{L \times d}$  with noisy input  $\varphi_t(x_0) = (1 - t)x_0 + tw_{r \rightarrow D^{1:L}}$  ( $t \sim \mathcal{U}[0, 1]$  and  $x_0 \sim \mathcal{N}(0^{1:L}, I)$ ).

For smooth transitions of sequences longer than the window length  $L$ , we incorporate last  $L'$  audio features and motion latents  $w_{r \rightarrow D^{-L':0}}$  from the preceding window as additional input.

The flow matching objective  $\mathcal{L}_{\text{OT}}(\theta)$  is defined by

$$\begin{aligned} \mathcal{L}_{\text{OT}}(\theta) = & \|v_t^{1:L}(x_t, \mathbf{c}_t; \theta) - u_t(x|w_{r \rightarrow D^{1:L}})\|, \\ & + \|v_t^{-L':0}(x_t, \mathbf{c}_t; \theta) - w_{r \rightarrow D^{-L':0}}\|, \end{aligned} \quad (11)$$

where  $x_t := [w_{r \rightarrow D^{-L':0}} | \varphi_t(x_0)] \in \mathbb{R}^{(-L'+L) \times d}$  is the concatenated input,  $\mathbf{c}_t \in \mathbb{R}^{(-L'+L) \times h}$  is the driving condition consisting of  $[t, w_{r \rightarrow S}, w_e, a^{1:L}, a^{-L':0}]$ . Note that  $w_e$  and  $w_{r \rightarrow S}$  are shared across the  $L' + L$  frames. We incorporate a velocity loss [67] to supervise temporal consistency:

$$\mathcal{L}_{\text{vel}}(\theta) = \|\Delta v_t - \Delta u_t\|, \quad (12)$$

where  $\Delta v_t$  and  $\Delta u_t$  are the one-frame difference along the time-axis for the prediction  $v_t \in \mathbb{R}^{(-L'+L) \times d}$  and the target  $[w_{r \rightarrow D^{-L':0}} | u_t] \in \mathbb{R}^{(-L'+L) \times d}$ , respectively.

The total objective  $\mathcal{L}_{\text{total}}(\theta)$  is

$$\mathcal{L}_{\text{total}}(\theta) = \lambda_{\text{OT}} \mathcal{L}_{\text{OT}}(\theta) + \lambda_{\text{vel}} \mathcal{L}_{\text{vel}}(\theta), \quad (13)$$

where  $\lambda_{\text{OT}}$  and  $\lambda_{\text{vel}}$  are the balancing coefficients. During training, we apply dropout to  $w_r$ ,  $w_e$ , and  $a^{1:L}$  with a probability of 0.1 for CFV. Additionally, we apply dropout to the preceding audio and motion latents with a probability 0.5 for smooth transition in the initial window.

**Inference** During inference, we sample the generating vector field from noise  $x_0$ , using the driving conditions  $w_{r \rightarrow S}$ ,  $w_e$ , and  $a^{1:L}$ , as well as the  $L'$  frames of preceding audio and generated motion latents.

We extend the CFV [10] to an incremental CFV to separately adjust the audio and emotion, inspired by [3]:

$$\begin{aligned} \tilde{v}_t \approx & v_t(x_0, \mathbf{c}_t |_{\{a^{1:L}, w_e\}}) \\ & + \gamma_a [v_t(x_0, \mathbf{c}_t |_{w_e}) - v_t(x_0, \mathbf{c}_t |_{\{a^{1:L}, w_e\}})] \\ & + \gamma_e [v_t(x_0, \mathbf{c}_t) - v_t(x_0, \mathbf{c}_t |_{w_e})], \end{aligned} \quad (14)$$

where  $\gamma_a$  and  $\gamma_e$  are the guidance scales for audio and emotion, respectively.  $\mathbf{c}_t |_{\{x, y\}}$  denotes the driving condition

Table 1. Quantitative comparison results with state-of-the-art methods on HDTF [85] / RAVDESS [39]. The best result for each metric is in **bold**, and the second-best result is underlined.  
<sup>†</sup>: evaluated with raw  $256 \times 256$  resolution outputs.

| Method                      | Image & Video Generation      |                                 |                             |                             |                             | Lip Synchronization         |                             |
|-----------------------------|-------------------------------|---------------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|
|                             | FID ↓                         | FVD ↓                           | CSIM ↑                      | E-FID ↓                     | P-FID ↓                     | LSE-D ↓                     | LSE-C ↑                     |
| SadTalker <sup>†</sup> [84] | 71.952 / 119.430              | 339.058 / 376.294               | 0.644 / 0.644               | 1.914 / 3.500               | 1.456 / 2.045               | 7.947 / <u>7.273</u>        | 7.305 / 4.748               |
| EDTalk <sup>†</sup> [66]    | 50.078 / 75.020               | 211.284 / 304.933               | 0.626 / 0.676               | 1.579 / 3.468               | 0.054 / 0.090               | 8.123 / 7.682               | 7.623 / <u>5.318</u>        |
| AniTalker <sup>†</sup> [36] | 39.512 / 70.430               | <u>184.454</u> / 265.341        | 0.643 / 0.725               | 1.830 / 2.330               | 0.092 / 0.126               | 7.907 / 8.176               | 7.288 / 4.555               |
| Hallo [79]                  | 25.363 / 57.648               | 197.196 / 375.557               | <b>0.869</b> / <b>0.860</b> | <b>1.039</b> / 2.492        | 0.037 / 0.050               | <u>7.792</u> / 7.613        | <u>7.582</u> / 4.795        |
| EchoMimic [7]               | 33.552 / 81.839               | 296.757 / 320.220               | 0.823 / 0.805               | 1.234 / 3.201               | <b>0.023</b> / <u>0.047</u> | 8.903 / 8.161               | 6.242 / 4.144               |
| <b>FLOAT (Ours)</b>         | <b>21.100</b> / <b>31.681</b> | <b>162.052</b> / <b>166.359</b> | <u>0.843</u> / <u>0.810</u> | <u>1.229</u> / <b>1.367</b> | <u>0.032</u> / <b>0.031</b> | <b>7.290</b> / <b>6.994</b> | <b>8.222</b> / <b>5.730</b> |




Figure 5. Qualitative comparison results with state-of-the-art methods on HDTF [85] / RAVDESS [39]. Please refer to supplementary videos. Note that we additionally provide a video comparison with EMO [68] and VASA-1 [80] using their video demonstration.

without the condition  $x$  and  $y$ . We set  $\gamma_a = 2$  and  $\gamma_e = 1$  based on the ablation studies on  $\gamma_a$  and  $\gamma_e$  provided in supplementary materials.

After sampling, ODE solver receives the estimated vector field to compute the motion latents through numerical integration. We empirically find that FLOAT can generate reasonable motion with around 10 number of function evaluations (NFE). Please refer to supplementary videos.

Lastly, we add the source identity latent to the generated motion latents and decode them into video frames using the motion latent decoder.

## 5. Experiments

### 5.1. Dataset and Pre-processing

For training the motion latent auto-encoder, we use three open-source datasets: **HDTF** [85], **RAVDESS** [39], and **VFHQ** [78]. When training FLOAT, we exclude VFHQ because it does not support the synchronized audio. HDTF [85] is for high-definition talking face generation, containing videos of over 300 unique identities. RAVDESS [39] includes more than 2,400 emotion-intensive videos of 24 different identities. VFHQ [78] is designed for high-resolution video super-resolution and includes a large num-

ber of unique identities, which compensates the limited number of identities of the preceding datasets. Following the strategy of [58], we first convert each video to 25 FPS and resample the audio into 16 kHz. Then, we crop and resize the facial region to  $512^2$  resolution [5]. After the pre-processing, for HDTF, we use a total of 11.3 hours of 240 videos featuring 230 different identities for training, and videos of 78 different identities, each 15 seconds long, for test. For RAVDESS, we use videos of 22 identities for training, and videos of the remaining 2 identities for test, with each 3-4 seconds long and representing 14 emotional intensities. Note that the identities in the training and test are disjoint in both datasets.

## 5.2. Implementation Details

The motion latent dimension is set to  $d = 512$  with  $M = 20$  distinct orthogonal directions. For the vector predictor, we use 8 attention heads, a hidden dimension  $h = 1024$ , and an attention window length  $T = 2$ . Considering the length of the training video clips, we set  $L = 50$  frames with preceding  $L' = 10$  frames at once, encompassing 2.4 seconds of video. We employ the Adam optimizer [30] with a batch size of 8 and a learning rate of  $10^{-5}$ . We use  $L1$  distance for the norm  $\|\cdot\|$  in the training objective. We set the balancing coefficients to  $\lambda_{OT} = \lambda_{vel} = 1$ . The entire training takes about 2 days for 2,000k steps on a single NVIDIA A100 GPU. We use Euler method [35] for the ODE solver.

## 5.3. Evaluation

**Metrics and Baselines** For evaluating the image and video generation quality, we measure Fréchet Inception Distance (FID) [57] and 16 frames Fréchet Video Distance (FVD) [69]. For facial identity, expression and head motion, we measure Cosine Similarity of identity embedding (CSIM) [11], Expression FID (E-FID) [68] and Pose FID (P-FID), respectively. Lastly, we measure Lip-Sync Error Distance and Confidence (LSE-D and LSE-C [51]) for audio-visual alignment.

We compare our method with state-of-the-art audio-driven talking portrait methods whose official implementations are publicly available. For non-diffusion methods, we compare with **SadTalker** [84] and **EDTalk** [66]. For diffusion methods, we compare with **AniTalker** [36], **Hallo** [79], and **EchoMimic** [7].

**Comparison Results** In Tab. 1 and Fig. 5, we show the quantitative and qualitative comparison results, respectively. FLOAT outperforms other methods on most of the metrics and visual quality in both datasets.

Additionally, we provide video comparison results with **EMO** [68] and **VASA-1** [80] in the supplementary materials, using their demonstration videos due to the infeasibility of direct implementation.



Figure 6. Test-time pose editing using  $\lambda$ -control ( $\lambda_{15}(\hat{D}) \pm 10$ ).

## 5.4. Applications

**Test-time Pose Editing via Orthonormal Basis  $V$**  Since FLOAT learns the underlying motion latent structure, it is natural to assume that for any *sampled* motion latent  $w_{r \rightarrow \hat{D}}$ , there exist motion coefficients  $\{\lambda_m(\hat{D})\}_{m=1}^M$  satisfying the representation in Eq. (9):  $w_{r \rightarrow \hat{D}} = \sum_{m=1}^M \lambda_m(\hat{D}) \cdot \mathbf{v}_m$ .

We can always compute these coefficients in *closed form* by taking inner products between the sampled motion  $w_{r \rightarrow \hat{D}}$  and the learned orthonormal basis  $V$ :

$$\langle w_{r \rightarrow \hat{D}}, \mathbf{v}_k \rangle = \langle \sum_{m=1}^M \lambda_m(\hat{D}) \cdot \mathbf{v}_m, \mathbf{v}_k \rangle = \lambda_k(\hat{D}), \quad (15)$$

where  $\langle \mathbf{v}_m, \mathbf{v}_k \rangle = \delta_{m,k}$  and  $\delta$  is Kronecker delta. At this point, we can edit the sampled motions by editing the corresponding coefficients (e.g., via linear operation) and combining them back into the motion latent. As shown in Fig. 6, it allows us to control head direction without interfering with other motions due to the orthogonality of the basis. We refer to this test-time editing technique as  $\lambda$ -control.

**Additional Driving Signals** In Fig. 7 and Tab. 2, we experiment with additional driving conditions, *head poses* and *image-driven emotion labels*, to explore additional controllability in our method. We employ 3DMM head pose parameters  $p \in \mathbb{R}^6$  [2] extracted by [12]. We concatenate a sequence of pose parameters  $p^{1:L} \in \mathbb{R}^{L \times 6}$  with the other driving conditions, and then map them to  $c_t^{1:L} \in \mathbb{R}^{L \times h}$ . We also experiment on image-driven emotion [56] for frame-wise emotion control rather than the long-term emotion enhancement. FLOAT can effectively accommodate these additional conditions, highlighting its flexibility across diverse control signals.

**Redirecting Speech-driven Emotion** Since FLOAT learns diverse emotions in the emotion-intensive data distribution [39], the generated emotion-aware motion can be modified by *redirecting* the speech-driven emotion label toward a different emotion at inference time. As illustrated in Fig. 8, this technique is particularly beneficial for manual redirection when the emotion predicted from speech is complex or ambiguous.




Figure 7. Additional conditioning results of FLOAT. *3DPose*, *S2E*, and *I2E* denote 3D head pose parameters [12], speech-to-emotion [49], and image-to-emotion [56], respectively.

Table 2. Quantitative results of FLOAT with additional conditions (HDTF [85] / RAVDSS [39]). *S2E*, *I2E*, and *3DPose* denote speech-to-emotion [49], image-to-emotion [56], and 3DMM pose parameters [12], respectively.

| Configurations        | FID ↓           | FVD ↓             | E-FID ↓       | P-FID ↓       | LSE-D ↓       |
|-----------------------|-----------------|-------------------|---------------|---------------|---------------|
| A <b>FLOAT (Ours)</b> | 21.100 / 31.681 | 162.052 / 166.359 | 1.229 / 1.367 | 0.032 / 0.031 | 7.290 / 6.994 |
| B A + 3DPose          | 19.721 / 29.721 | 126.663 / 112.894 | 0.926 / 1.152 | 0.012 / 0.016 | 7.516 / 7.047 |
| C A - S2E             | 21.235 / 32.035 | 155.032 / 166.866 | 1.254 / 1.502 | 0.031 / 0.025 | 7.264 / 7.222 |
| D A - S2E + I2E       | 21.528 / 31.609 | 158.577 / 162.369 | 1.158 / 1.305 | 0.034 / 0.022 | 7.183 / 7.150 |



Figure 8. Redirecting the unclear emotion prediction to a desirable one-hot encoding, which can be further intensified by the CFV.

## 5.5. Ablation Studies

**Ablation on Frame-wise AdaLN** We compare frame-wise AdaLN (and gating) followed by masked self-attention to separate conditioning from attending, with a cross-attention that performs conditioning and attending simultaneously. As shown in Tab. 3, both approaches achieve competitive image and video quality, while frame-wise AdaLN provides better expression generation and lip synchronization. We observe that frame-wise AdaLN can achieve more diverse head motions than the cross-attention. Please refer to supplementary videos.

**Ablation on Flow Matching** We compare flow matching with two types of diffusion models:  $\epsilon$ -prediction (noise) and  $x_0$ -prediction (signal) [52, 67]. In both cases, we adopt our vector predictor architecture as denoising networks. We adopt diffusion training settings of VASA-1 [80] (500 diffusion steps with a cosine noise scheduler [46] and 50 DDIM denoising steps) for the indirect comparison with [80]. Notably, diffusion and flow matching achieve competitive results on image quality while the latter achieves the better lip synchronization. In Fig. 9, we compare the forward pass efficiency by measuring frames per second (FPS) of each

Table 3. Ablation studies of FLOAT on HDTF [85]. The best result for each metric is in **bold**, and the second-best result is underlined.

| Method                             | FID ↓         | FVD ↓          | E-FID ↓      | LSE-D ↓      | # NFEs ↓  |
|------------------------------------|---------------|----------------|--------------|--------------|-----------|
| Ours (w. Cross-Attn.)              | 21.873        | 162.702        | 1.452        | <u>7.757</u> | <b>10</b> |
| Ours (w. Diff., $\epsilon$ -pred.) | <u>21.190</u> | <b>161.666</b> | <b>1.213</b> | 9.922        | 50        |
| Ours (w. Diff., $x_0$ -pred.)      | 21.697        | 162.847        | 1.278        | 9.048        | 50        |
| <b>FLOAT (Ours)</b>                | <b>21.100</b> | <u>162.052</u> | <u>1.229</u> | <b>7.290</b> | <b>10</b> |




Figure 9. Comparison of the forward pass efficiency. We compute FPS on a single NVIDIA V100 GPU.

model. Thanks to the compact motion latent representation and OT-based flow matching, FLOAT achieves the highest FPS, superior lip-sync performance, dynamic head motion, and the lowest NFEs.

## 6. Conclusion

We proposed FLOAT, a flow matching based audio-driven talking portrait generation model leveraging a learned motion latent space. We introduced a transformer-based vector field predictor, enabling temporally consistent motion generation. Additionally, we incorporated speech-driven emotion labels into the motion sampling process to improve the naturalness of the audio-driven talking motions. FLOAT addresses current core limitations of diffusion-based talking portrait video generation methods by reducing the sampling time through flow matching while achieving the remarkable sample quality. Extensive experiments verified that FLOAT achieves state-of-the-art performance in terms of visual quality, motion fidelity, and efficiency.

**Discussion** We leave further discussion considering *limitations*, *future work*, and *ethical considerations* in the supplementary materials.

## References

- [1] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems*, 33:12449–12460, 2020. 5
- [2] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194, 1999. 3, 7
- [3] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18392–18402, 2023. 5
- [4] Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, et al. Video generation models as world simulators, 2024. 2
- [5] Adrian Bulat and Georgios Tzimiropoulos. How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230,000 3d facial landmarks). In *International Conference on Computer Vision*, 2017. 7
- [6] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018. 2, 3
- [7] Zhiyuan Chen, Jiajiong Cao, Zhiquan Chen, Yuming Li, and Chenguang Ma. Echomimic: Lifelike audio-driven portrait animations through editable landmark conditions. *arXiv preprint arXiv:2407.08136*, 2024. 1, 2, 3, 4, 6, 7
- [8] Kun Cheng, Xiaodong Cun, Yong Zhang, Menghan Xia, Fei Yin, Mingrui Zhu, Xuan Wang, Jue Wang, and Nannan Wang. Videoretalking: Audio-based lip synchronization for talking head video editing in the wild. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022. 1, 2
- [9] Joon Son Chung and Andrew Zisserman. Out of time: automated lip sync in the wild. In *Asian Conference on Computer Vision*, pages 251–263, 2016. 1, 2
- [10] Quan Dao, Hao Phung, Binh Nguyen, and Anh Tran. Flow matching in latent space. *arXiv preprint arXiv:2307.08698*, 2023. 2, 3, 5
- [11] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4690–4699, 2019. 7
- [12] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019. 7, 8
- [13] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. 2
- [14] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016. 2
- [15] Nikita Drobyshev, Jenya Chelishev, Taras Khakhulin, Aleksei Ivakhnenko, Victor Lempitsky, and Egor Zakharov. Megaportraits: One-shot megapixel neural head avatars. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 2663–2671, 2022. 3
- [16] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024. 2
- [17] Johannes S Fischer, Ming Gui, Pingchuan Ma, Nick Stracke, Stefan A Baumann, and Björn Ommer. Boosting latent diffusion with flow matching. *arXiv preprint arXiv:2312.07360*, 2023. 2
- [18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 2
- [19] Jiazhi Guan, Zhanwang Zhang, Hang Zhou, Tianshu Hu, Kaisiyuan Wang, Dongliang He, Haocheng Feng, Jingtuo Liu, Errui Ding, Ziwei Liu, et al. Stylesync: High-fidelity generalized and personalized lip sync in style-based generator. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1505–1515, 2023. 1, 2
- [20] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. *arXiv preprint arXiv:2307.04725*, 2023. 2, 3
- [21] Tianyu He, Junliang Guo, Runyi Yu, Yuchi Wang, Jialiang Zhu, Kaikai An, Leyi Li, Xu Tan, Chunyu Wang, Han Hu, et al. Gaia: Zero-shot talking avatar generation. *arXiv preprint arXiv:2311.15230*, 2023. 1, 2
- [22] Geoffrey Hinton. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 2
- [23] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2
- [24] Li Hu. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8153–8163, 2024. 2, 3, 4
- [25] 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*, pages 1–10, 2022. 2, 3, 5
- [26] Jianwen Jiang, Chao Liang, Jiaqi Yang, Gaojie Lin, Tianyun Zhong, and Yanbo Zheng. Loopy: Taming audio-driven portrait avatar with long-term motion dependency. *arXiv preprint arXiv:2409.02634*, 2024. 1, 2, 3
- [27] Minguk Kang, Richard Zhang, Connelly Barnes, Sylvain Paris, Suha Kwak, Jaesik Park, Eli Shechtman, Jun-Yan Zhu,

- and Taesung Park. Distilling diffusion models into conditional gans. *arXiv preprint arXiv:2405.05967*, 2024. 2
- [28] Taekyung Ki and Dongchan Min. Stylelipsync: Style-based personalized lip-sync video generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22841–22850, 2023. 1, 2
- [29] Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 1, 2
- [30] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 7
- [31] Tobias Kirschstein, Simon Giebenhain, and Matthias Nießner. Diffusionavatars: Deferred diffusion for high-fidelity 3d head avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5481–5492, 2024. 2
- [32] Matthew Le, Apoorv Vyas, Bowen Shi, Brian Karrer, Leda Sari, Rashel Moritz, Mary Williamson, Vimal Manohar, Yossi Adi, Jay Mahadeokar, et al. Voicebox: Text-guided multilingual universal speech generation at scale. *Advances in neural information processing systems*, 36, 2024. 2
- [33] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *ArXiv e-prints*, pages arXiv–1607, 2016. 5
- [34] Senmao Li, Taihang Hu, Fahad Shahbaz Khan, Linxuan Li, Shiqi Yang, Yaxing Wang, Ming-Ming Cheng, and Jian Yang. Faster diffusion: Rethinking the role of unet encoder in diffusion models. *arXiv preprint arXiv:2312.09608*, 2023. 2
- [35] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022. 2, 3, 4, 7
- [36] Tao Liu, Feilong Chen, Shuai Fan, Chenpeng Du, Qi Chen, Xie Chen, and Kai Yu. Anitalker: Animate vivid and diverse talking faces through identity-decoupled facial motion encoding. *arXiv preprint arXiv:2405.03121*, 2024. 1, 3, 6, 7
- [37] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022. 2, 4
- [38] Xingchao Liu, Xiwen Zhang, Jianzhu Ma, Jian Peng, et al. Instaflow: One step is enough for high-quality diffusion-based text-to-image generation. In *The Twelfth International Conference on Learning Representations*, 2023. 2
- [39] Steven R Livingstone and Frank A Russo. The ryerson audio-visual database of emotional speech and song (ravdess): A dynamic, multimodal set of facial and vocal expressions in north american english. *PloS one*, 13(5): e0196391, 2018. 6, 7, 8
- [40] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35:5775–5787, 2022. 2
- [41] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *arXiv preprint arXiv:2211.01095*, 2022. 2
- [42] Simian Luo, Yiqin Tan, Longbo Huang, Jian Li, and Hang Zhao. Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378*, 2023. 2
- [43] Yifeng Ma, Suzhen Wang, Zhipeng Hu, Changjie Fan, Tangjie Lv, Yu Ding, Zhidong Deng, and Xin Yu. Styletalk: One-shot talking head generation with controllable speaking styles. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1896–1904, 2023. 2
- [44] Yifeng Ma, Shiwei Zhang, Jiayu Wang, Xiang Wang, Yingya Zhang, and Zhidong Deng. Dreamtalk: When expressive talking head generation meets diffusion probabilistic models. *arXiv preprint arXiv:2312.09767*, 2023. 2, 3
- [45] Dongchan Min, Minyoung Song, Eunji Ko, and Sung Ju Hwang. Styletalker: One-shot style-based audio-driven talking head video generation. *arXiv preprint arXiv:2208.10922*, 2022. 1, 2
- [46] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pages 8162–8171. PMLR, 2021. 2, 8
- [47] Se Jin Park, Minsu Kim, Joanna Hong, Jeongsoo Choi, and Yong Man Ro. Synctalkface: Talking face generation with precise lip-syncing via audio-lip memory. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2062–2070, 2022. 1
- [48] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023. 2, 4
- [49] Leonardo Pepino, Pablo Riera, and Luciana Ferrer. Emotion recognition from speech using wav2vec 2.0 embeddings. *arXiv preprint arXiv:2104.03502*, 2021. 5, 8
- [50] Adam Polyak, Amit Zohar, Andrew Brown, Andros Tjandra, Animesh Sinha, Ann Lee, Apoorv Vyas, Bowen Shi, Chih-Yao Ma, Ching-Yao Chuang, et al. Movie gen: A cast of media foundation models. *arXiv preprint arXiv:2410.13720*, 2024. 2
- [51] KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Namboodiri, and CV Jawahar. A lip sync expert is all you need for speech to lip generation in the wild. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 484–492, 2020. 1, 2, 7
- [52] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1 (2):3, 2022. 2, 8
- [53] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR, 2015. 1, 2
- [54] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, 2022. 2, 3, 4
- [55] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour,

- Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35:36479–36494, 2022. 2
- [56] Andrey V Savchenko. Hsemotion: High-speed emotion recognition library. *Software Impacts*, 14:100433, 2022. 5, 7, 8
- [57] Maximilian Seitzer. pytorch-fid: FID Score for PyTorch. <https://github.com/mseitzer/pytorch-fid>, 2020. Version 0.3.0. 7
- [58] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. *Advances in neural information processing systems*, 32, 2019. 7
- [59] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 4
- [60] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2
- [61] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 2
- [62] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *arXiv preprint arXiv:2303.01469*, 2023. 2
- [63] Michał Stypułkowski, Konstantinos Vougioukas, Sen He, Maciej Zięba, Stavros Petridis, and Maja Pantic. Diffused heads: Diffusion models beat gans on talking-face generation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5091–5100, 2024. 2
- [64] Zhiyao Sun, Tian Lv, Sheng Ye, Matthieu Lin, Jenny Sheng, Yu-Hui Wen, Mingjing Yu, and Yong-jin Liu. Diffposetalk: Speech-driven stylistic 3d facial animation and head pose generation via diffusion models. *ACM Transactions on Graphics (TOG)*, 43(4):1–9, 2024. 2
- [65] Shuai Tan, Bin Ji, and Ye Pan. Emmn: Emotional motion memory network for audio-driven emotional talking face generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22146–22156, 2023. 2, 3
- [66] Shuai Tan, Bin Ji, Mengxiao Bi, and Ye Pan. Edtalk: Efficient disentanglement for emotional talking head synthesis. In *European Conference on Computer Vision*, pages 398–416. Springer, 2025. 1, 6, 7
- [67] 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. 2, 5, 8
- [68] Linrui Tian, Qi Wang, Bang Zhang, and Liefeng Bo. Emo: Emote portrait alive-generating expressive portrait videos with audio2video diffusion model under weak conditions. *arXiv preprint arXiv:2402.17485*, 2024. 2, 3, 4, 6, 7
- [69] Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018. 7
- [70] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. 2, 4, 5
- [71] Cong Wang, Kuan Tian, Jun Zhang, Yonghang Guan, Feng Luo, Fei Shen, Zhiwei Jiang, Qing Gu, Xiao Han, and Wei Yang. V-express: Conditional dropout for progressive training of portrait video generation. *arXiv preprint arXiv:2406.02511*, 2024. 2, 3
- [72] Kaisiyuan Wang, Qianyi Wu, Linsen Song, Zhuoqian Yang, Wayne Wu, Chen Qian, Ran He, Yu Qiao, and Chen Change Loy. Mead: A large-scale audio-visual dataset for emotional talking-face generation. In *European Conference on Computer Vision*, pages 700–717. Springer, 2020. 3, 5
- [73] Suzhen Wang, Lincheng Li, Yu Ding, Changjie Fan, and Xin Yu. Audio2head: Audio-driven one-shot talking-head generation with natural head motion. *arXiv preprint arXiv:2107.09293*, 2021. 2
- [74] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10039–10049, 2021. 3
- [75] Yaohui Wang, Di Yang, Francois Bremond, and Antitza Dantcheva. Latent image animator: Learning to animate images via latent space navigation. *arXiv preprint arXiv:2203.09043*, 2022. 2, 4
- [76] Huawei Wei, Zejun Yang, and Zhisheng Wang. Aniportrait: Audio-driven synthesis of photorealistic portrait animation. *arXiv preprint arXiv:2403.17694*, 2024. 2, 3
- [77] Yibo Xia, Lizhen Wang, Xiang Deng, Xiaoyan Luo, and Yebin Liu. Gmtalker: Gaussian mixture based emotional talking video portraits. *arXiv preprint arXiv:2312.07669*, 2023. 3
- [78] Liangbin Xie, Xintao Wang, Honglun Zhang, Chao Dong, and Ying Shan. Vfhq: A high-quality dataset and benchmark for video face super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 657–666, 2022. 6
- [79] Mingwang Xu, Hui Li, Qingkun Su, Hanlin Shang, Liwei Zhang, Ce Liu, Jingdong Wang, Luc Van Gool, Yao Yao, and Siyu Zhu. Hallo: Hierarchical audio-driven visual synthesis for portrait image animation. *arXiv preprint arXiv:2406.08801*, 2024. 2, 3, 4, 6, 7
- [80] Sicheng Xu, Guojun Chen, Yu-Xiao Guo, Jiaolong Yang, Chong Li, Zhenyu Zang, Yizhong Zhang, Xin Tong, and Baining Guo. Vasa-1: Lifelike audio-driven talking faces generated in real time. *arXiv preprint arXiv:2404.10667*, 2024. 2, 3, 5, 6, 7, 8
- [81] Fei Yin, Yong Zhang, Xiaodong Cun, Mingdeng Cao, Yanbo Fan, Xuan Wang, Qingyan Bai, Baoyuan Wu, Jue Wang, and Yujiu Yang. Styleheat: One-shot high-resolution editable talking face generation via pre-trained stylegan. In *European conference on computer vision*, pages 85–101. Springer, 2022. 2
- [82] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. 2

- [83] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 586–595, 2018. [4](#)
- [84] Wenxuan Zhang, Xiaodong Cun, Xuan Wang, Yong Zhang, Xi Shen, Yu Guo, Ying Shan, and Fei Wang. Sadtalker: Learning realistic 3d motion coefficients for stylized audio-driven single image talking face animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8652–8661, 2023. [1](#), [2](#), [6](#), [7](#)
- [85] Zhimeng Zhang, Lincheng Li, Yu Ding, and Changjie Fan. Flow-guided one-shot talking face generation with a high-resolution audio-visual dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3661–3670, 2021. [6](#), [8](#)
- [86] Zhimeng Zhang, Zhipeng Hu, Wenjin Deng, Changjie Fan, Tangjie Lv, and Yu Ding. Dinet: Deformation inpainting network for realistic face visually dubbing on high resolution video. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3543–3551, 2023. [1](#)
- [87] Hang Zhou, Yasheng Sun, Wayne Wu, Chen Change Loy, Xiaogang Wang, and Ziwei Liu. Pose-controllable talking face generation by implicitly modularized audio-visual representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4176–4186, 2021. [2](#)
- [88] Yang Zhou, Xintong Han, Eli Shechtman, Jose Echevarria, Evangelos Kalogerakis, and Dingzeyu Li. Makelttalk: speaker-aware talking-head animation. *ACM Transactions On Graphics (TOG)*, 39(6):1–15, 2020. [2](#)
- [89] Shenhao Zhu, Junming Leo Chen, Zuozhuo Dai, Yinghui Xu, Xun Cao, Yao Yao, Hao Zhu, and Siyu Zhu. Champ: Controllable and consistent human image animation with 3d parametric guidance. *arXiv preprint arXiv:2403.14781*, 2024. [4](#)