Ditto: Motion-Space Diffusion for Controllable Realtime Talking Head Synthesis

Tianqi Li,
Ant Group
shijian.ltq@antgroup.com

Ruobing Zheng†,
Ant Group
zhengruobing.zrb@antgroup.com

Minghui Yang
Ant Group
minghui.ymh@antgroup.com

Jingdong Chen
Ant Group
jingdongchen.cjd@antgroup.com

Ming Yang
Ant Group
m.yang@antgroup.com

Project Page: https://digital-avatar.github.io/ai/Ditto/

Abstract

Recent advances in diffusion models have endowed talking head synthesis with subtle expressions and vivid head movements, but have also led to slow inference speed and insufficient control over generated results. To address these issues, we propose Ditto, a diffusion-based talking head framework that enables fine-grained controls and real-time inference. Specifically, we utilize an off-the-shelf motion extractor and devise a diffusion transformer to generate representations in a specific motion space. We optimize the model architecture and training strategy to address the issues in generating motion representations, including insufficient disentanglement between motion and identity, and large internal discrepancies within the representation. Besides, we employ diverse conditional signals while establishing a mapping between motion representation and facial semantics, enabling control over the generation process and correction of the results. Moreover, we jointly optimize the holistic framework to enable streaming processing, real-time inference, and low first-frame delay, offering functionalities crucial for interactive applications such as AI assistants. Extensive experimental results demonstrate that Ditto generates compelling talking head videos and exhibits superiority in both controllability and real-time performance.

1 Introduction

Audio-driven talking head synthesis [1, 2, 3, 4, 5, 6, 7] has emerged as a prominent research topic in recent years. In particular, one-shot methods [8, 9, 10, 11] substantially reduce the data prerequisite for target identities, no longer relying on templates. Early approaches based on Generative Adversarial Networks (GANs) [12] are the predominant solutions, showing accurate lip synchronization but still lacking diversity and realism. Recently, diffusion-based techniques [13, 14, 15], with EMO [3] as a notable example, have made substantial progress. These methods maintain lip synchronization and capture subtle facial expressions, synchronizing head movements with speech rhythm, which are crucial traits to achieve compelling synthesized results.

However, two critical issues hinder the broad application of these diffusion-based methods. (1) lack of control over facial motion. Existing approaches struggle to enforce fine-grained control over talking heads, including facial movements, basic emotions, and head rotations. This leaves users with no direct means to adjust the results in addition to regeneration. As the generation quality is relatively random, achieving desired results becomes challenging and potentially time-consuming. (2) Slow inference speed. Currently, most methods can hardly achieve real-time inference on a single GPU, which is a critical requirement for interactive scenarios, such as AI assistants and live video streaming. VASA-1 [16] innovatively achieves real-time inference through a two-stage generation approach, demonstrat-ing the feasibility of training a Diffusion Transformer (DiT) [17] based on a motion space. However, due to its source code is not publicly available, related follow-up work is limited. On the other hand, VASA-1 uses an implicit motion representation, which does not support control or adjustment of the generated results.

In this paper, we propose Ditto, a diffusion-based talking head framework that enables fine-grained controls and real-time inference. Specifically, we utilize an off-the-shelf motion extractor [18] and devise a conditional diffusion transformer to generate representations in a specific motion space. We optimize the model architecture and training strategy to address inherent issues in motion representations, including insufficient disentanglement between motion and identity information, and large internal discrepancies within the motion representations.

Besides, we employ diverse conditional signals to achieve fine-grained control over the generation process. We also establish a mapping between motion representation and facial semantics, enabling fine-tuning of facial movements and correction of visual defects.

Moreover, we jointly optimize the holistic framework to enable streaming processing, real-time inference, and low first-frame delay, offering functionalities crucial for interactive applications such as AI assistants. Extensive experimental results demonstrate that Ditto generates compelling talking head videos and exhibits superiority in both controllability and real-time performance. For the advancement of the community, we make our source code available for open-source use.
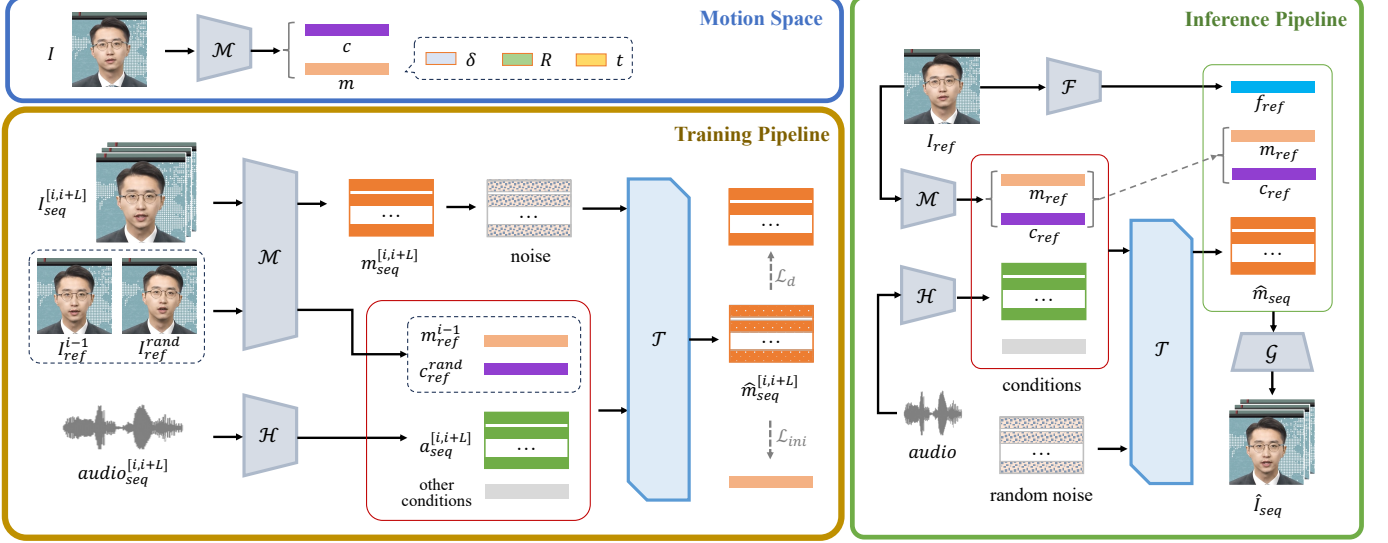
---

† Corresponding Author

Figure 1: The summary of the proposed **Ditto**, where $\mathcal{T}$ is the DiT for motion generation, $\mathcal{H}$ is the HuBERT for audio feature extraction, $\mathcal{M}$ is the Motion Extractor, $\mathcal{F}$ is the Appearance Feature Extractor and $\mathcal{G}$ is the Face Renderer.

2 RELATED WORK

In audio-driven talking head synthesis, one-shot methods have garnered research attention due to their ability to minimize data requirements to a single photograph. While early GAN-based approaches have succeeded in generating photorealistic textures and accurate lip movements, they fall short in capturing realistic expressions and head movements. Recent diffusion-based methods have made substantial progress in addressing these limitations, truly achieving the generation of realistic talking-head videos without relying on carefully curated templates. EMO [3] stands out as a groundbreaking approach that leverages audio and other "weak controls" to train a conditional diffusion model for end-to-end video generation, enhancing the vividness and realism of the generated results. Subsequent methods have widely adopted its innovative Dual U-Net architecture and two-stage training scheme. Notable follow-up works include EchoMimic [14], Hallo [15], and Loopy [13], which have improved the original framework in various aspects such as hierarchical audio-visual mapping, diverse driving signals, and long-term motion modeling, respectively. Nevertheless, these methods still rely on denoising and rendering images in a general VAE space. This redundant, implicit latent space not only increases the learning complexity of the Diffusion model but also results in slow inference speeds. VASA-1 [16] introduces a novel approach by training a Diffusion Transformer (DiT) [17] within a motion-appearance disentangled space, defined by a face reenactment model [19]. This latent space, tailored specifically for the facial domain, combined with a two-stage generation process, significantly reduces inference time while maintaining vibrant results.

3 METHOD

In this section, we introduce the detailed implementation of the Ditto framework. We begin with a brief introduction to the motion space. Then we discuss the model architecture and training strategy to address inherent issues in motion represen-

tations. Finally, we introduce a novel motion control approach and demonstrate how to correct visual defects using gaze as an example.

3.1 Motion Space

While Latent Diffusion Models (LDMs) [20] significantly reduce computational requirements by operating in latent space rather than pixel space, their pretrained VAE-defined latent representations present two key limitations in talking head synthesis. (1) The latent space remains redundant, with entangled motion and texture, introducing overhead in both diffusion training and pixel-level rendering. (2) These representations are implicit and lack explicit correspondence with facial attributes.

Due to the shared underlying facial geometry, individuals across different identities exhibit highly consistent facial motion patterns. Recent facial reenactment works [21, 22, 23, 24, 25, 26] also fully utilize this characteristic. We construct the motion space based on LivePortrait [18], one of the state-of-the-art methods in this field. We aim to let the diffusion model generate universal motion, and then incorporate identity information during pixel-level rendering.

Specifically, a single frame I is processed through Motion Extractor $\mathcal{M}$ derived from [18] to yield canonical keypoints $\mathbf{c} \in \mathbb{R}^{K \times 3}$, expression deformations $\boldsymbol{\delta} \in \mathbb{R}^{K \times 3}$, and headposes $\mathbf{R} \in \mathbb{R}^{3 \times 3}$ with translations $\mathbf{t} \in \mathbb{R}^3$. The $\mathbf{m} = \{\boldsymbol{\delta}, \mathbf{R}, \mathbf{t}\}$ serve as the identity-agnostic motion representation, which is used for training the diffusion model to predict the corresponding audio-driven $\hat{\mathbf{m}} = \{\hat{\boldsymbol{\delta}}, \hat{\mathbf{R}}, \hat{\mathbf{t}}\}$, as illustrated in Figure 1.

Then a one-shot face renderer, comprising an Appearance Feature Extractor $\mathcal{F}$ and a Face Renderer $\mathcal{G}$ [18], is used to synthesize talking head videos. Given a target portrait, $\mathcal{M}$ first extracts its canonical keypoints $\mathbf{c}_{\mathbf{ref}}$ and motion representation $\mathbf{m}_{\mathbf{ref}} = \{\boldsymbol{\delta}_{\mathbf{ref}}, \mathbf{R}_{\mathbf{ref}}, \mathbf{t}_{\mathbf{ref}}\}$. Then the reference and generated im-

plicit 3D keypoints ($\mathbf{x_{ref}}$ and $\hat{\mathbf{x}}$) are computed as:

$$\mathbf{x_{ref}} = \mathbf{c_{ref}}\mathbf{R_{ref}} + \delta_{ref} + \mathbf{t_{ref}}, \quad (1)$$

$$\hat{\mathbf{x}} = \mathbf{c_{ref}}\hat{\mathbf{R}} + \hat{\delta} + \hat{\mathbf{t}}. \quad (2)$$

These two sets of implicit keypoints guide $\mathcal{G}$ in learning a warping field that transforms the appearance features $\mathbf{f_{ref}}$ from $\mathcal{F}$ to fit the generated motions and decodes the warped features into image space to synthesize frame $\hat{I}$ as:

$$\hat{I} = \mathcal{G}(\mathbf{f_{ref}}, \mathbf{x_{ref}}, \hat{\mathbf{x}}). \quad (3)$$

3.2 Motion Generation with Diffusion Transformer

As shown in Figure 2, we employ a Conditional Diffusion Transformer (DiT) for audio-to-motion generation. Along with the **audio feature a**, we incorporate multiple auxiliary conditional signals as follows.

3.2.1 Conditional Signals

Considering the fact that even LivePortrait [18] is hard to fully disentangle motion and identity information, and the impure motion representations increase the difficulty of the training process. We use the canonical keypoints as the **identity feature $\mathbf{c_{ref}}$** to guide the generation of facial motions, ensuring the generated motions match the target facial characteristics and can appropriately animate the target identity. This identity-adaptive setting mitigates the long-standing issue of incomplete disentanglement between motion and identity [27, 28]. Ablation experiments also confirm the above conclusions.

We also extract emotions [29] from video frames and employ a unique **emotion label s** at the clip level, establishing a direct and controllable correspondence instead of simply learning the audio-to-expression mapping. **Eye state e** is composed of the aspect ratio and relative pupil position [30], governing blinking and gaze which have weak correlations with audio. A **reference initial motion $\mathbf{m_{ref}}$** serves as a guide for the motion generation within each clip, which can be utilized to enhance inter-clip motion continuity, as well as to periodically steer the motion towards a specified target state, thereby mitigating error accumulation in long sequence generation [31].

For **e**, $\mathbf{c_{ref}}$, and **s**, we use them as the Enhanced Conditional Signals (ECS) which are aligned temporally with audio features **a**, followed by channel-wise concatenation before feeding into a cross-attention module to guide the generation process. As for $\mathbf{m_{ref}}$, we use it as an Initial Conditional Signal (ICS) by replicating and concatenating it channel-wise with the noise sequence, aiming to influence motion generation primarily in the initial stages. In summary, we denote our input condition as $\mathbf{C} = \{\mathbf{a}, \mathbf{e}, \mathbf{c_{ref}}, \mathbf{s}, \mathbf{m_{ref}}\}$

3.2.2 Training Strategy

We devise specialized strategies for training the DiT to accommodate the specific motion representations.

Horizontal Flip: Facial motions associated with speech are predominantly symmetrical. However, in-the-wild talking-head training data often exhibits an uneven distribution of head orientations. This imbalance can affect the extracted motion features,
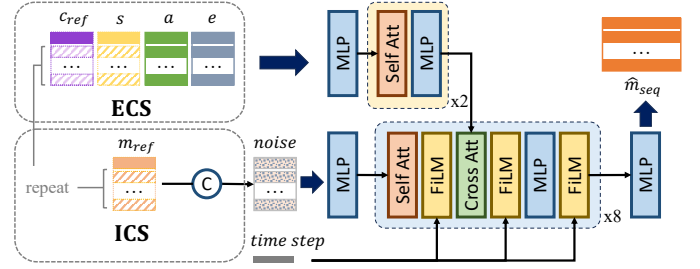

Figure 2: Visualization of the architecture of the proposed conditional Diffusion Transformer, which generates compact motion representations based on various conditional signals.

potentially biasing generated motions towards one side. To address this issue, we horizontally flipped the face images, to augment the training data to a balanced audio-to-motion correspondence for both sides.

Adaptive Loss Weights: Different components of the motion representation exhibit variations in their movement patterns. For example, motions such as lip movements, facial expressions, eye blinking, and head pose differ significantly in their relationships with driving signals like audio, and in the magnitude of their movements. It is difficult to train the DiT to weigh all motion representations uniformly. Therefore, we group the motion representations based on their control regions, and dynamically adjust the weights of each group throughout the training process. We calculate the difference in average loss between the current epoch and the previous epoch, then apply softmax normalization across groups to derive adjustment coefficients for the loss weights. Experiments demonstrate that this dynamic weighting strategy improves both the loss convergence and the quality of generated motion.

Validation Metric: Given that lip movements are the primary descriptor of facial motion, we employ a lipsync score [1] as our validation metric. We render the generated motions from the validation videos and use this metric to assist in selecting the preferred model checkpoint. This approach addresses a common training issue in diffusion models, where the loss curve often fails to reflect the generation quality accurately.

Loss Function: Following DDPM [32], a denoising network $\mathcal{T}$ is trained by minimizing the mean square error loss as follows:

$$\mathcal{L}_d = \mathbb{E}_{t\sim\mathcal{U}[1,T],\mathbf{m_0},\mathbf{C}} \left[\|\mathbf{m_0} - \mathcal{T}(\mathbf{m}_t, \mathbf{C}, t)\|_2^2 \right], \quad (4)$$

where $\mathbf{C}$ is the conditional signal, $\mathbf{m}_t$ is a perturbed version of real motion data $\mathbf{m_0}$ by adding t step noises. To enhance temporal stability, we also regress the velocities and accelerations of motion movements.

$$\mathcal{L}_t = \left\| \hat{\mathbf{m}}' - \mathbf{m}' \right\|_2^2 + \left\| \hat{\mathbf{m}}'' - \mathbf{m}'' \right\|_2^2, \quad (5)$$

where $\mathbf{m}''$ and $\mathbf{m}'$ denote the first-order and second-order derivatives of motion representations $\mathbf{m}$. Furthermore, to enhance the initial motion guidance, we compute an initial motion loss $\mathcal{L}_{ini}$ between the reference initial motion and the first element of the generated motion clip. In summary, the final loss can be expressed as:

$$\mathcal{L} = \mathcal{L}_d + \mathcal{L}_t + \mathcal{L}_{ini} \quad (6)$$
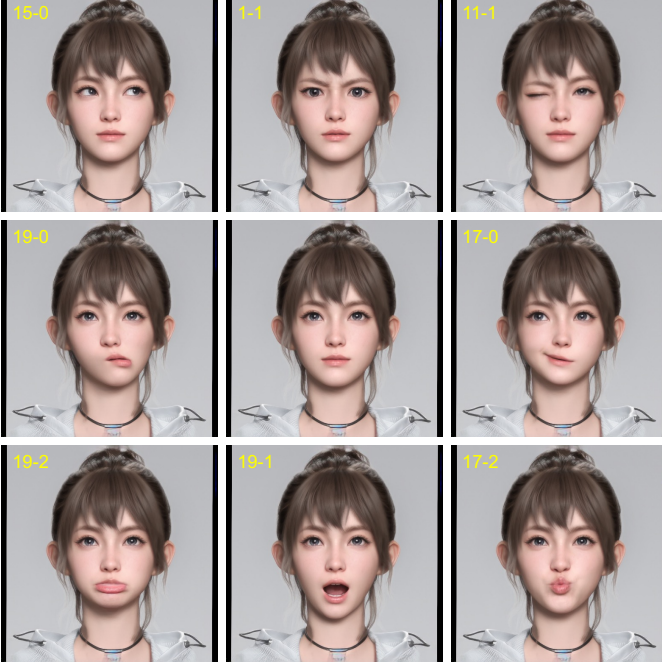
Figure 3: Examples of dedicated fine-grained control of subtle facial motion. The notation "15-0" in the top-left image denotes the application of an x-axis offset to the 15th keypoint. The central image depicts the rendered result when applying zero offsets across all dimensions.

3.2.3 Motion Control

Unlike headpose which can be directly controlled via rotation and translation, our implicit keypoint-based motion representations [21, 22, 23] lack clear correspondence with facial attributes. To address this limitation, [18] employs landmark-guided optimization to align implicit keypoints with explicit landmarks. We extend this work by establishing the direct mapping between deformations and facial semantics, enabling motion control in a blendshape-like manner. Specifically, we represent expression deformation as a 63-D vector, corresponding to the x, y, and z coordinates of 21 implicit 3D keypoints. For each dimension, we apply small positive and negative offsets and render facial images for a fixed identity. As illustrated in Figure 3, each dimension of deformations affects relatively distinct facial regions. For instance, the 34th dimension governs the right eye's opening and closing motion. The 58th dimension controls mouth open, analogous to the "jaw open" blendshape in ARKit. Therefore, we enable two types of generation control: (1) Regional control, where motion generation can be constrained to specific local facial areas. For example, limiting movements to facial features close to the neck for better integration with the torso, and (2) Magnitude control, where we can impose constraints on deformation values to prevent artifacts such as unnatural facial expressions.

Next, we introduce how to correct visual defects in the generated motions through a case of gaze adjustment. In the early test, we discover that the generated gaze is bound to the head pose. This results in the avatar's gaze constantly wandering with head movements in the generated video, unable to focus on the camera. The reason is that we control the gaze by adding the per frame eye state $\mathbf{e}$ during training. But we use the fixed eye state during inference, which is extracted from the single reference image.

To address this issue, we record a template video, requiring the actor to smoothly rotate head in various angles while maintaining eye focus on the camera and a constant facial expression. Then we extract the corresponding head poses $\mathbf{R}$ and expression deformations δ. By subtracting the initial frame's $\mathbf{R_0}$ and δ_0 from the $\mathbf{R}$ and δ of each frame, we obtain synchronized variation data for pose and motion. Since the actor's expression remains constant throughout the video, the motion variation only reflects the change in gaze. We use a regression $\delta_e = \mathcal{K}(\mathbf{R_e})$ to obtain the mapping between the gaze variation δ_e and the head pose variation $\mathbf{R_e}$, which is used to correct the generated gaze based on the current headpose, $\delta_{correct} = \delta + \delta_e$.

3.3 Realtime Streaming Inference

Realtime dialogue scenarios impose two critical requirements on talking head methods: (1) low-latency inference within each module and (2) efficient data streaming between modules. To cope with these requirements, we optimize the inference process across three primary modules: audio feature extraction, motion generation, and video synthesis.

3.3.1 Audio Feature

We employ the HuBERT [33] for audio feature extraction, incorporating transformer [34] acceleration techniques to enable realtime streaming audio processing. Specifically, we adopt a KV cache strategy [35] to provide pseudo-contextual information extracted from a fixed corpus, concatenated with incoming audio stream segments. This approach allows the HuBERT model pre-trained on long utterances to maintain high-quality features even on extremely short audio segments. Complementarily, we implement causal mask to ensure proper attention allocation to the actual input while significantly reducing computational complexity. These optimizations enable our model to process 0.4-second audio stream units in realtime on CPU environment. We utilize these enhanced features for both model training and inference.

3.3.2 Motion Generation

We align the audio features with the video frame rate and segment the streaming audio features into fixed-length intervals, incorporating a predefined overlap between adjacent segments. We leverage the long video generation strategy introduced in [36], where we substitute progressive latent fusion with segment-wise fusion on generated motion sequences to support streaming output. During fusion, weights are determined by the relative position of each frame within its segment, with frames closer to the segment center receiving higher weights.

To achieve realtime performance, we reduce the number of denoising steps in DiT inference from 50 to 10. Benefiting from the motion representation, the inference results at 10 and 50 steps maintain comparable quality in evaluation. The DiT model is converted to TensorRT for execution on GPU environment.

Table 1: **Comparisons with existing methods on the Talk9 dataset**. The s* represents the number of denoising steps. The **best** and <u>second best</u> results are in **bold** and <u>underline</u> specifically.

Method	FID ↓	FVD ↓	CSIM ↑	Sync-C ↑	Sync-D ↓	RTF ↓
GT	-	-	-	8.044	6.943	-
MuseTalk	21.445	436.862	0.807	5.586	8.400	2.248
EchoMimic	42.554	395.754	0.840	5.733	9.204	35.528
Hallo	22.996	271.680	0.812	7.652	7.590	53.082
Hallo2	22.899	245.236	0.806	7.737	7.608	56.838
Ours-s50	<u>17.254</u>	**219.368**	**0.864**	<u>8.069</u>	**7.114**	<u>2.121</u>
Ours-s10	**17.060**	<u>231.182</u>	<u>0.861</u>	**8.111**	<u>7.291</u>	**0.635**

Table 2: The quantitative comparisons on the HDTF100 dataset. *Data with † are from Hallo2 paper.*

Method	FID ↓	FVD ↓	SyncC ↑	SyncD ↓
EchoMimic [†]	47.331	532.733	5.930	9.143
Hallo [†]	16.748	366.066	7.268	7.714
Hallo2 [†]	<u>16.616</u>	<u>239.517</u>	<u>7.379</u>	<u>7.697</u>
Ours	**16.430**	**134.640**	**8.939**	**6.913**

Table 3: User study results.

Method	Visual quality	Lipsync	Naturalness
EchoMimic	54.0%	38.0%	58.7%
Hallo	21.3%	36.0%	33.3%
Hallo2	40.7%	45.3%	**59.3%**
Ours	**84.0%**	**80.7%**	48.7%

3.3.3 Video Synthesis

Our one-shot face renderer sequentially generates photorealistic images from motion sequences. Features relevant to the target identity are pre-extracted and stored. The TensorRT-optimized renderer enables realtime inference and video stream output on GPUs. If video output is required, a CPU-based FFmpeg can parallelly compress the video stream.

4 EXPERIMENTS

4.1 Implementation Details

For training data, we collect and clean about 50 hours of broadcast scene videos from 330 identities, with an average video length of 150 seconds. For evaluation, due to the limited computing resources, we conduct our primary quantitative comparison on *Talk9* dataset while providing additional quantitative comparison on *HDTF100* dataset. The *Talk9* dataset consists of 9 videos from [37, 38, 39, 40] and has been widely used in GAN/NeRF-based talking head evaluations. *HDTF100* contains 100 randomly sampled videos from the HDTF [41] dataset, which has been widely used for quantitative evaluation in previous works [14, 15, 42]. To evaluate the generalizability of our method, we select diverse image and audio samples from [15, 14] for qualitative evaluations

In experiments, the headpose euler angles are represented using discrete bins [21], and the final motion representation contains 265 dimensions. We have trained our models using 8 NVIDIA A100 GPUs with a batch size of 1024, using the Adan [43] optimizer with a learning rate of 1e-4 and 0.02 weight decay for 500 epochs. We select the preferred checkpoint based on the lipsync score.

During training, audio features, video frames, and corresponding conditional signals are aligned at 25 fps. We sample segments of L=80-frames (3.2s) from video clips to train the Diffusion

Transformer. The reference initial motion $\mathbf{m}_{\mathbf{ref}}$ is derived from the preceding frame of the current training segment, and the reference canonical keypoints $\mathbf{c}_{\mathbf{ref}}$ are taken from a random frame of the same video clip. During inference, both the $\mathbf{m}_{\mathbf{ref}}$ and $\mathbf{c}_{\mathbf{ref}}$ are derived from the source image.

4.2 Baseline and Metrics

We compare our method with recent publicly available implementations, including the non-diffusion-based lip editing method MuseTalk [44], and the diffusion-based talking head generation methods EchoMimic [14], Hallo [15] and Hallo2 [42]. The following metrics are used to evaluate the generated results: Fréchet Inception Distance (FID) [45] for frame-level quality, Fréchet Video Distance (FVD) [46] for video-level quality, Cosine Similarity (CSIM) of face recognition features [47] for identity consistency, Sync-C and Sync-D [1] for audio-visual synchronization. We also evaluate the inference performance by Real-Time Factor (RTF) and First-Frame Delay (FFD). RTF < 1 and an FFD as low as possible are the basic requirements for realtime streaming inference.

4.3 Quantitative Evaluation

The quantitative results on Talk9 and HDTF100 datasets are shown in Table 1 and Table 2, respectively. Our method achieves the best results among competing methods on all evaluated metrics. The motion representation facilitates the learning of the diffusion model and makes the lipsync score close to the real video, which is the most important criterion for evaluating talking head tasks. How to maintain identity consistency, measured by the CSIM score, is a long-standing challenge for one-shot methods. Our method achieves higher CSIM scores than other methods, attributed to the disentanglement of appearance and motion. Our method also effectively improves the frame quality and inter-frame continuity of the generated video obtaining the lowest FID and FVD among all methods. Table 1 also reports the results of reducing the number of denoising steps. The results of

Figure 4: Generation results with portraits of different styles and scales. Fine-grained control over gaze, emotion, pose, etc.
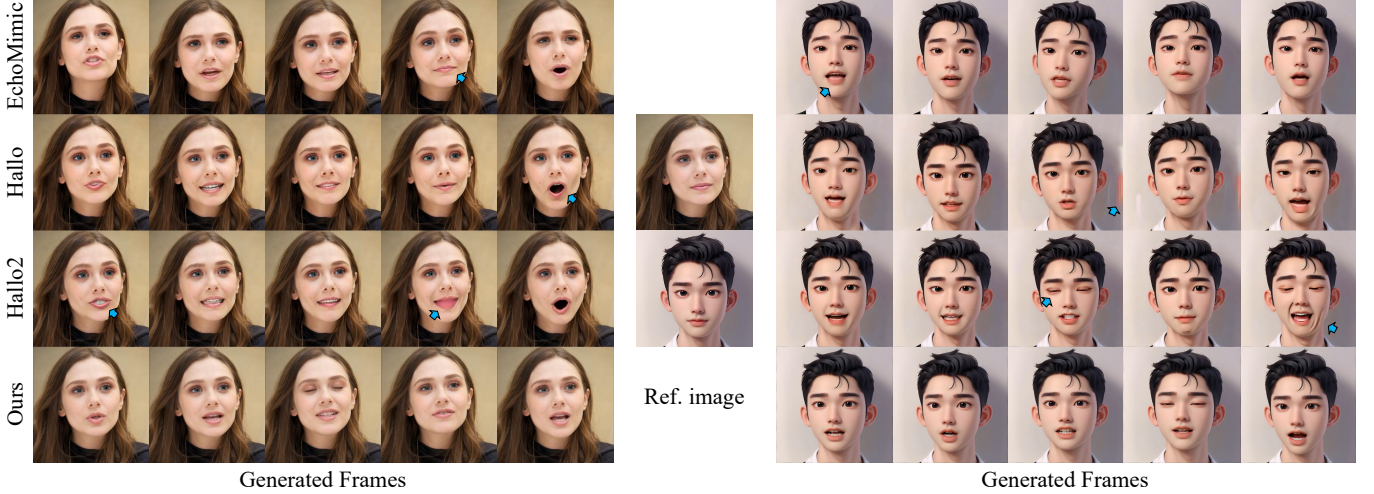
Figure 5: The qualitative comparison of our approach on two characters with different styles, poses, and hairstyles. For each character, we generate videos using the same audio input with each method and select frames at the same location for comparison. Blue arrows indicate the locations of artifacts, including inaccurate lip movements, disordered teeth, blur, etc.

Table 4: Ablation study results. **C-kp** represents canonical keypoints condition, **Emo** represents emotion label condition, and **Ada-w** represents Adaptive Loss Weights.

Method	FID ↓	FVD ↓	CSIM ↑	SyncC ↑	SyncD ↓
Full Model	17.254	219.368	0.864	8.069	7.114
w/o C-kp	17.438	261.693	0.849	6.932	8.476
w/o Emo	23.036	324.078	0.792	6.310	8.990
w/o Ada-w	18.978	588.821	0.844	0.289	14.992

10-step and 50-step denoising exhibit minimal disparity in end-to-end evaluation, substantiating the efficacy of our inference optimization. Analysis of the RTF reveals that diffusion-based methods exhibit inference speeds 30-50 times slower than realtime. In contrast, our proposed method achieves inference capabilities that surpass realtime performance, and even outpace the non-diffusion-based lip-editing method MuseTalk.

User Study. To estimate the quality of our method and SOTAs from human perspectives, we conduct a blind user study with 10 participants. Specifically, we randomly select 5 image-audio pairs to generate 5 driving clips for each tested method, resulting in a total of 20 clips. Each participant is presented with two videos generated by different methods for the same set of inputs, and asked to choose the better one in terms of visual quality, lip synchronization, and naturalness of facial expressions and head movements. This process is repeated C_4^2 times. The results are summarized in Table 3, where our method outperforms other methods in visual quality and lip synchronization. We also notice that Ditto performs slightly weaker in the measure of "naturalness of facial expressions and head movement". We believe this is due to the motion representation attenuating the learning of some high-frequency motion information, which still requires our continuous optimization.

4.4 Qualitative Evaluation

The qualitative results are presented in Figures 4 and Figure 5. Figure 5 shows a qualitative comparison of our approach on two characters with different styles, poses, and hairstyles. For each character, we generate videos using the same audio input with each method and select frames at the same location for comparison. Our method shows advantages in the consistency of generated details. For example, the shape and texture of teeth generated by baselines change between frames. The baselines also have problems such as generating exaggerated expressions, blur, and inaccurate lip movements. In addition, the generation results of Hallo and Hallo2 for stylized portraits may deteriorate gradually over time, such as excessive facial wrinkles, abnormal shadows, and background artifacts. In contrast, our method shows stable performance across diverse portrait styles.

Figure 4 demonstrates that Ditto can not only handle portraits of different styles but also support fine-grained generation control, such as gaze, pose, emotion, etc. Furthermore, our approach enables seamless stitching of the generated head with the original body, facilitating full-body portrait animation without the misalignment issues often associated with end-to-end diffusion methods. This makes Ditto readily integrate with the body-driven model, achieving strong generation capabilities in a wide range of applications.

4.5 Ablation Study

Conditional Information. We evaluate the impact of different controlling signals on the end-to-end generation results. As shown in Table 4, when canonical keypoints are excluded (**w/o C-kp**), all metrics drop, especially for FVD and Sync. This suggests that the identity-related geometry provided by canonical keypoints enhances the learning process of diffusion, particularly in adapting generated motion to the target identity. Conversely, imperfections in motion generation directly affect dynamic visual quality and lip synchronization accuracy.
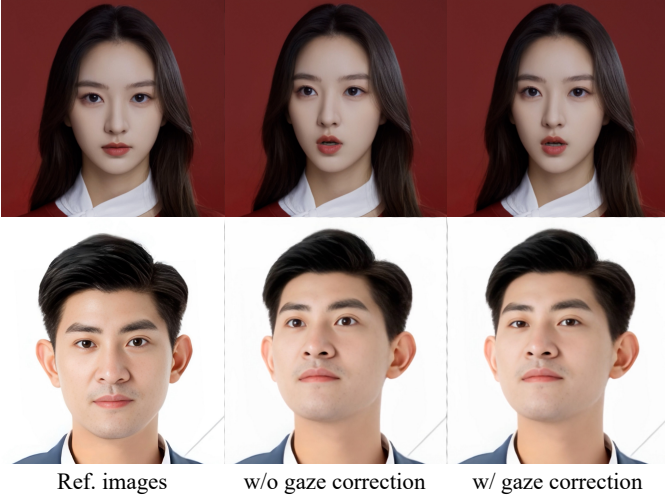
Figure 6: The gaze adjustment decouples gaze direction from the head pose, enabling the avatar's eyes to focus on the camera, producing more natural and engaging eye contact.

Table 5: The time consumption of single-step inference in different modules and corresponding real-time factor (RTF).

Module	Audio2Feat	Motion DiT	Face Rendering
inference	23ms	62ms	15ms
RTF	0.115	0.310	0.375

When emotion labels are omitted (**w/o Emo**), we observe a significant decline across all metrics. Without direct emotional control, learning facial expressions solely from audio becomes more challenging, resulting in a decrease in video quality and lip synchronization. In the supplementary materials, we demonstrate the consequent issues such as abrupt expression changes and exaggerated facial expressions, which also impact identity consistency throughout the video.

Adaptive Loss Weights. Without adaptive loss weights (**w/o Ada-w**), the problem of unbalanced convergence between pose and expression becomes severe, hindering the model to learn the correspondence between audio and facial motion. We observe that the generated videos exhibit minimal facial motions, resulting in poor Sync and FVD scores, while CSIM and FID metrics are less affected.

Gaze Adjustment. As shown in Figure 6, without gaze adjustment, the gaze direction is rigidly bound to the head pose. The gaze adjustment decouples gaze direction from the head pose, enabling the avatar's eyes to focus on the camera, producing more natural and engaging eye contact.

4.6 Inference Performance

We separately evaluate the inference performance of the proposed method from both modular and end-to-end perspectives. We evaluate module-wise performance using both inference time and valid Real-Time Factor (RTF), which represents the single-step inference time divided by the valid segment length (excluding the overlapping part). The test environment is a 12-

Table 6: Inference performance for offline video output and online streaming output at the head region (512×512) and full-body scales (1920×1080).

Pipeline	RTF	FFD
talking-head offline	0.635	-
talking-head online	0.895	385ms
full-body offline	0.648	-
full-body online	0.914	392ms

core Intel(R) Xeon(R) Platinum 8369B CPU @ 2.90GHz, 1 NVIDIA A100 GPU, and 100G memory.

Table 5 shows that each module's RTF is under 1, guaranteeing realtime end-to-end inference for the pipeline. In Table 6, we present the inference performance for offline video output and online streaming output at two scales: the head region and full-body portrait. Compared to offline video output, online streaming output requires longer inter-segment overlap to ensure low First-Frame Delay (FFD), resulting in a higher RTF. Full-body output is only marginally slower than head-only output, primarily due to the process of integrating the head into the body. Our method ensures end-to-end inference with RTF below 1 and FFD under 400ms, meeting the requirements for realtime interactive scenarios.

5 CONCLUSION

We propose Ditto, a diffusion-based talking head framework that enables fine-grained controls and real-time inference. We devise a Conditional DiT and corresponding training strategy to address the issues in generating motion representations. We jointly optimize the holistic framework to enable streaming processing, real-time inference, and low first-frame delay. For the advancement of the community, we make our source code available for open-source use.

REFERENCES

[1] KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Namboodiri, and CV Jawahar. A lip sync expert is all you need for speech to lip generation in the wild. In *Proceedings of the 28th ACM international conference on multimedia*, pages 484–492, 2020.

[2] Nikita Drobyshev, Antoni Bigata Casademunt, Konstantinos Vougioukas, Zoe Landgraf, Stavros Petridis, and Maja Pantic. Emoportraits: Emotion-enhanced multimodal one-shot head avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8498–8507, 2024.

[3] 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.

[4] Zicheng Zhang, Ruobing Zheng, Bonan Li, Congying Han, Tianqi Li, Meng Wang, Tiande Guo, Jingdong Chen, Ziwen Liu, and Ming Yang. Learning dynamic tetrahedra for high-quality talking head synthesis. In *Proceedings of the*

IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 5209–5219, 2024.

[5] Ruobing Zheng, Bo Song, and Changjiang Ji. Learning pose-adaptive lip sync with cascaded temporal convolutional network. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4255–4259. IEEE, 2021.

[6] 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. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 6696–6705, 2024.

[7] Enric Corona, Andrei Zanfir, Eduard Gabriel Bazavan, Nikos Kolotouros, Thiemo Alldieck, and Cristian Sminchisescu. Vlogger: Multimodal diffusion for embodied avatar synthesis. *arXiv preprint arXiv:2403.08764*, 2024.

[8] Weichuang Li, Longhao Zhang, Dong Wang, Bin Zhao, Zhigang Wang, Mulin Chen, Bang Zhang, Zhongjian Wang, Liefeng Bo, and Xuelong Li. One-shot high-fidelity talking-head synthesis with deformable neural radiance field. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17969–17978, 2023.

[9] 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*, pages 8652–8661, 2023.

[10] Huawei Wei, Zejun Yang, and Zhisheng Wang. Aniportrait: Audio-driven synthesis of photorealistic portrait animation. *arXiv preprint arXiv:2403.17694*, 2024.

[11] Zhenhui Ye, Tianyun Zhong, Yi Ren, Jiaqi Yang, Weichuang Li, Jiawei Huang, Ziyue Jiang, Jinzheng He, Rongjie Huang, Jinglin Liu, et al. Real3d-portrait: One-shot realistic 3d talking portrait synthesis. *arXiv preprint arXiv:2401.08503*, 2024.

[12] 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.

[13] 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.

[14] 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.

[15] 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.

[16] 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.

[17] 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.

[18] Jianzhu Guo, Dingyun Zhang, Xiaoqiang Liu, Zhizhou Zhong, Yuan Zhang, Pengfei Wan, and Di Zhang. Liveportrait: Efficient portrait animation with stitching and retargeting control. *arXiv preprint arXiv:2407.03168*, 2024.

[19] 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.

[20] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Bjorn Ommer. High-resolution image synthesis with latent diffusion models. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun 2022. doi: 10.1109/cvpr52688.2022.01042. URL http://dx.doi.org/10.1109/cvpr52688.2022.01042.

[21] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun 2021. doi: 10.1109/cvpr46437.2021.00991. URL http://dx.doi.org/10.1109/cvpr46437.2021.00991.

[22] Fa-Ting Hong, Longhao Zhang, Li Shen, and Dan Xu. Depth-aware generative adversarial network for talking head video generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3397–3406, 2022.

[23] 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.

[24] Aliaksandr Siarohin, Oliver J Woodford, Jian Ren, Menglei Chai, and Sergey Tulyakov. Motion representations for articulated animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13653–13662, 2021.

[25] Jian Zhao and Hui Zhang. Thin-plate spline motion model for image animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3657–3666, 2022.

[26] Bowen Zhang, Chenyang Qi, Pan Zhang, Bo Zhang, HsiangTao Wu, Dong Chen, Qifeng Chen, Yong Wang, and Fang Wen. Metaportrait: Identity-preserving talking head generation with fast personalized adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22096–22105, 2023.

[27] Yu Deng, Duomin Wang, Xiaohang Ren, Xingyu Chen, and Baoyuan Wang. Portrait4d: Learning one-shot 4d head avatar synthesis using synthetic data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7119–7130, 2024.

[28] Phong Tran, Egor Zakharov, Long-Nhat Ho, Anh Tuan Tran, Liwen Hu, and Hao Li. Voodoo 3d: Volumetric portrait disentanglement for one-shot 3d head reenactment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10336–10348, 2024.

[29] Andrey V Savchenko. Hsemotion: High-speed emotion recognition library. *Software Impacts*, 14:100433, 2022.

[30] Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Guang Yong, Juhyun Lee, et al. Mediapipe: A framework for building perception pipelines. *arXiv preprint arXiv:1906.08172*, 2019.

[31] Xiang Wang, Shiwei Zhang, Changxin Gao, Jiayu Wang, Xiaoqiang Zhou, Yingya Zhang, Luxin Yan, and Nong Sang. Unianimate: Taming unified video diffusion models for consistent human image animation. *arXiv preprint arXiv:2406.01188*, 2024.

[32] Jonathan Ho, Ajay Jain, Pieter Abbeel, and UC Berkeley. Denoising diffusion probabilistic models. *arXiv preprint arXiv:2006.11239*, 2020. URL https://arxiv.org/abs/2006.11239.

[33] Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM transactions on audio, speech, and language processing*, 29:3451–3460, 2021.

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

[35] Zihang Dai. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*, 2019.

[36] Yuang Zhang, Jiaxi Gu, Li-Wen Wang, Han Wang, Junqi Cheng, Yuefeng Zhu, and Fangyuan Zou. Mimicmotion: High-quality human motion video generation with confidence-aware pose guidance. *arXiv preprint arXiv:2406.19680*, 2024.

[37] Yudong Guo, Keyu Chen, Sen Liang, Yong-Jin Liu, Hujun Bao, and Juyong Zhang. Ad-nerf: Audio driven neural radiance fields for talking head synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5784–5794, 2021.

[38] Yuanxun Lu, Jinxiang Chai, and Xun Cao. Live speech portraits: Real-time photorealistic talking-head animation. *ACM Trans. Graph.*, 40(6), dec 2021. ISSN 0730-0301. doi: 10.1145/3478513.3480484. URL https://doi.org/10.1145/3478513.3480484.

[39] 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 *European Conference on Computer Vision*, pages 666–682. Springer, 2022.

[40] Jiahe Li, Jiawei Zhang, Xiao Bai, Jun Zhou, and Lin Gu. Efficient region-aware neural radiance fields for high-fidelity talking portrait synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7568–7578, 2023.

[41] 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*, pages 3661–3670, 2021.

[42] Jiahao Cui, Hui Li, Yao Yao, Hao Zhu, Hanlin Shang, Kaihui Cheng, Hang Zhou, Siyu Zhu, and Jingdong Wang. Hallo2: Long-duration and high-resolution audio-driven portrait image animation. *arXiv preprint arXiv:2410.07718*, 2024.

[43] Xingyu Xie, Pan Zhou, Huan Li, Zhouchen Lin, and Shuicheng Yan. Adan: Adaptive nesterov momentum algorithm for faster optimizing deep models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.

[44] Yue Zhang, Minhao Liu, Zhaokang Chen, Bin Wu, Yubin Zeng, Chao Zhan, Yingjie He, Junxin Huang, and Wenjiang Zhou. Musetalk: Real-time high quality lip synchronization with latent space inpainting. *arXiv preprint arXiv:2410.10122*, 2024.

[45] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.

[46] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphaël Marinier, Marcin Michalski, and Sylvain Gelly. Fvd: A new metric for video generation. 2019.

[47] 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*, pages 4690–4699, 2019.