



# Follow-Your-Emoji: Fine-Controllable and Expressive Freestyle Portrait Animation

YUE MA, Hong Kong University of Science and Technology, Hong Kong  
 HONGYU LIU\*, Hong Kong University of Science and Technology, Hong Kong  
 HONGFA WANG\*, Tencent, China  
 HENG PAN\*, Tencent, China  
 YINGQING HE, Hong Kong University of Science and Technology, Hong Kong  
 JUNKUN YUAN, Tencent, China  
 AILING ZENG, Tencent, China  
 CHENGFEI CAI, Tencent, China  
 HEUNG-YEUNG SHUM, Tsinghua University, China  
 WEI LIU†, Tencent, China  
 QIFENG CHEN†, Hong Kong University of Science and Technology, Hong Kong

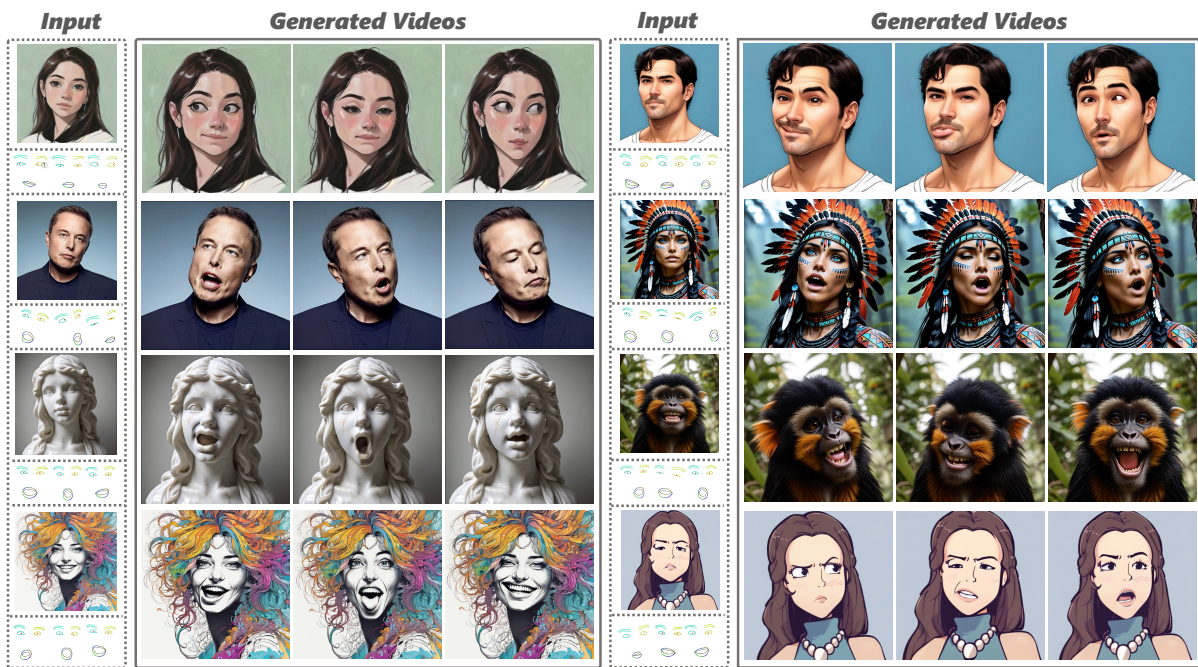


Fig. 1. **Qualitative results of our Follow-Your-Emoji.** The images of the input column are the reference portrait and the corresponding motion landmarks. Using exaggerated expressions with landmark sequences, our portrait animation framework can animate freestyle reference portraits, e.g., cartoons, realism, sculptures, and even animals. The input images are from Civitai [civ 2023].

\*Equal contribution.

†Joint corresponding authors.

Authors' Contact Information: Yue Ma, Hong Kong University of Science and Technology, Hong Kong, mayuefighting@gmail.com; Hongyu Liu, Hong Kong University of Science and Technology, Hong Kong, hliudq@cse.ust.hk; Hongfa Wang, Tencent, China, hongfawang@tencent.com; Heng Pan, Tencent, China, hengpan@tencent.com; Yingqing He, Hong Kong University of Science and Technology, Hong Kong, yhebm@connect.ust.hk; Junkun Yuan, Tencent, China, junkunyuan@tencent.com; Ailing Zeng, Tencent, China, ailingzengzzz@gmail.com; Chengfei Cai, Tencent, China, fletchercai@tencent.com; Heung-Yeung Shum, Tsinghua University, China, hshum@ust.hk; Wei Liu, Tencent, China, wl2223@columbia.edu; Qifeng Chen, Hong Kong University of Science and Technology, Hong Kong, cqf@ust.hk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed

We present Follow-Your-Emoji, a diffusion-based framework for portrait animation, which animates a reference portrait with target landmark sequences. The main challenge of portrait animation is to preserve the identity of the

for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SA Conference Papers '24, December 03–06, 2024, Tokyo, Japan

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1131-2/24/12

<https://doi.org/10.1145/3680528.3687587>

reference portrait and transfer the target expression to this portrait while maintaining temporal consistency and fidelity. To address these challenges, Follow-Your-Emoji equipped the powerful Stable Diffusion model with two well-designed technologies. Specifically, we first adopt a new explicit motion signal, namely expression-aware landmark, to guide the animation process. We discover this landmark can not only ensure the accurate motion alignment between the reference portrait and target motion during inference but also increase the ability to portray exaggerated expressions (i.e., large pupil movements) and avoid identity leakage. Then, we propose a facial fine-grained loss to improve the model’s ability of subtle expression perception and reference portrait appearance reconstruction by using both expression and facial masks. Accordingly, our method demonstrates significant performance in controlling the expression of freestyle portraits, including real humans, cartoons, sculptures, and even animals. By leveraging a simple and effective progressive generation strategy, we extend our model to stable long-term animation, thus increasing its potential application value. To address the lack of a benchmark for this field, we introduce **EmojiBench**, a comprehensive benchmark comprising diverse portrait images, driving videos, and landmarks. We show extensive evaluations on EmojiBench to verify the superiority of Follow-Your-Emoji. The code, training dataset and benchmark will be found in <https://github.com/mayuelala/FollowYourEmoji>.

CCS Concepts: • **Computing methodologies** → **Animation**.

Additional Key Words and Phrases: Portrait Animation, Diffusion Model, Expression-aware landmark, Facial Fine-Grained Loss

#### ACM Reference Format:

Yue Ma, Hongyu Liu, Hongfa Wang, Heng Pan, Yingqing He, Junkun Yuan, Ailing Zeng, Chengfei Cai, Heung-Yeung Shum, Wei Liu, and Qifeng Chen. 2024. Follow-Your-Emoji: Fine-Controllable and Expressive Freestyle Portrait Animation. In *SIGGRAPH Asia 2024 Conference Papers (SA Conference Papers '24)*, December 03–06, 2024, Tokyo, Japan. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3680528.3687587>

## 1 introduction

We study the task of portrait animation, which transfers the target sequences of poses and expressions from the driven video to the reference portrait. Combined with the generative adversarial network [Goodfellow et al. 2020] (GAN) and diffusion model [Sohl-Dickstein et al. 2015], recent portrait animation methods demonstrate widespread potential applications, such as online conferencing, virtual characters, and augmented reality.

For the GAN-based portrait animation method [Drobyshev et al. 2022; Liu et al. 2023a; Siarohin et al. 2019; Wang et al. 2021], they typically utilize a two-stage pipeline which first warps the reference image in feature space with flow field, then adopts the GAN as a rendering decoder to refine the warping features and generate the missing or occluded body parts. However, due to the limited performance of GAN and the inaccuracy of motion representation of the flow field, the generation results of these methods always suffer from unrealistic content and remarkable artifacts. In recent years, diffusion models [Ho et al. 2020; Song et al. 2020] have showcased better generation ability than GAN. Some methods bring powerful foundation diffusion models for high-quality video [Blattmann et al. 2023; Chen et al. 2023, 2024; Guo et al. 2023; He et al. 2022a; Ho et al. 2022a,b; Wu et al. 2023] and image generation [Ramesh et al. 2022; Rombach et al. 2021; Saharia et al. 2022] with large-scale image or video datasets. However, these foundation models can not directly handle the main challenges of the portrait animation task:

preserving the reference portrait’s identity during animation and effectively modeling the target expression for the portrait.

Intuitively, some methods [Chang et al. 2024; Hu et al. 2023; Wang et al. 2023a; Xu et al. 2024b; Zhu et al. 2024] try to modify the architecture of foundation diffusion model (i.e., Stable Diffusion [Rombach et al. 2021]) with some plug and play modules for portrait animation task and leverage the pretrained diffusion model as powerful prior information. Specifically, they utilize an appearance net [Hu et al. 2023] and CLIP model [Radford et al. 2021] to extract identity information of the reference portrait and temporal attention to establish temporal consistency between frames. However, the video results of these methods exhibit distortions and unrealistic artifacts, especially when animating uncommon domain portraits (i.e., cartoons, sculptures, and animals) that are not represented in the training data. We find this is mainly due to two reasons: (1) The motion representation (i.e., 2D landmarks [Chang et al. 2024; Hu et al. 2023] or the motion image itself [Xie et al. 2024]) adopted in these methods are not robust enough. During inference, 2D landmarks can easily lead to a misalignment between the facial features of the reference portrait and the target motion, resulting in identity leakage. However, setting the motion image itself as the signal needs to utilize third-party methods to change the identity of the target motion videos for training, as mentioned in Xportrait [Xie et al. 2024]. And it will destroy the subtle expression features in the original motion videos. (2) These methods utilize the original loss in the diffusion model during training, which is unsuitable for portrait animation tasks that need the model to focus on capturing reference facial appearance and expression changes.

In this paper, we present Follow-Your-Emoji, a novel diffusion-based framework for portrait animation. Apart from the commonly used appearance net and temporal attention in recent diffusion-based portrait animation methods, we propose several effective technologies to address the aforementioned problems. (1) We introduce the expression-aware landmark, a novel expression control signal, to guide the driving process more effectively. Specifically, we obtain the landmark by projecting the 3D keypoints obtained from MediaPipe [Lugaresi et al. 2019]. Owing to the inherent canonical property of 3D keypoints, we can effectively align the target motion with the reference portrait during inference, thereby avoiding identity leakage. However, MediaPipe is not robust enough, as the facial contour sometimes fails to conform to the face accurately. Consequently, the process of projecting landmarks has been modified to exclude facial contours and incorporate pupil points. This operation enables the model to better focus on expression changes (i.e., pupil point motion) while preventing it from influencing the shape and destroying the identity information of the reference portrait through the wrong facial contour. (2) We propose a facial fine-grained loss function to aid the model in focusing on capturing subtle expression changes and the detailed appearance of the reference portrait. Specifically, we first leverage both facial masks and expression masks with our expression-aware landmark, then compute the spatial distance between the ground truth and predicted results in these mask regions.

Through the aforementioned improvements, our approach can effectively drive freestyle portraits, as illustrated in Figure 1. Additionally, to train our model, we construct a high-quality expression

training dataset with 18 exaggerated expressions and 20-minute real-human videos from 115 subjects. We employ a progressive generation strategy that enables our method to scale to long-term animation synthesis with high fidelity and stability. To address the lack of a benchmark in portrait animation, we introduce a comprehensive benchmark called *EmojiBench*, which consists of 410 various style portrait animation videos that showcase a wide range of facial expressions and head poses. Finally, we conduct a comprehensive evaluation of *Follow-Your-Emoji* using *EmojiBench*. The evaluation results demonstrate the impressive performance of our method in handling portraits and motions that were outside of the training domain. Compared with the existing baseline methods, our method performs quantitatively and qualitatively better, delivering exceptional visual fidelity, faithful representation of identities, and precise motion rendering. In summary, our contributions can be summarized as follows:

- We introduce *Follow-Your-Emoji*, a diffusion-based framework for fine-controllable portrait animation. Based on the proposed progressive generation strategy, it can further produce long-term animation.
- To facilitate freestyle portrait animation, we propose the expression-aware landmarks as the motion representation and a facial fine-grained loss to help the diffusion model enhance the generation quality of facial expressions.
- To train our model, we introduce a new expression training dataset with 18 expressions and 20-min talking videos from 115 subjects. To validate the effectiveness of our methods, we construct a benchmark *EmojiBench*, and comprehensive results show the superiority of our *Follow-Your-Emoji* in fine-controllable and expressive aspects.

## 2 Related Work

### 2.1 Single-forward Portrait Animation

Animating a single portrait has attracted a lot of attention in the research. Previous approaches [Averbuch-Elor et al. 2017; Chang et al. 2024; Drobyshev et al. 2022; Kim et al. 2018; Siarohin et al. 2019; Thies et al. 2016; Wiles et al. 2018] mainly leverage Generative Adversarial Networks (GANs) [Goodfellow et al. 2020] to generate plausible motion using self-supervised learning. Due all GAN method methods can work without a discriminator, we consider them as the single-forward methods. The pioneering works primarily involved two steps: warping and rendering. These methods firstly estimate head and facial motion with open-source 2D/3D pose predictors [Lugaresi et al. 2019; Yang et al. 2023]. The facial representation is warped and fed into a generative model to synthesize dynamic frames with realistic animation and rich details. Following such a paradigm, a majority of approaches [Hong et al. 2022; Qu et al. 2023; Wang et al. 2021; Zhao and Zhang 2022] focus on improve facial warping estimation, including 3D neural landmarks [Wang et al. 2021], thin-plate splines [Zhao and Zhang 2022] and depth [Hong et al. 2022]. Additionally, the 3D morphable is utilized to model the expression and motion in *ReenacArtFace* [Qu et al. 2023]. *ToonTalker* [Gong

et al. 2023] employs the transformer architecture to help the warping process of cross-domain datasets. *MegaPortraits* [Drobyshev et al. 2022] enhances rendered image quality using high-resolution image data, whereas *FADM* [Zeng et al. 2023] enriches generated details using the proposed coarse-to-fine animation framework. *Face Vid2Vid* [Wang et al. 2021] presents a pure neural rendering to decompose identity-specific and motion-related information unsupervisedly. In addition to video reenactment, there are also various driving signals, such as 3D facial prior [Deng et al. 2020; Feng et al. 2021; Fried et al. 2019; Khakhulin et al. 2022; Sun et al. 2023; Xu et al. 2023] and audio [He et al. 2023; Tian et al. 2024; Xu et al. 2024a; Zhang et al. 2023a]. However, these methods primarily focus on talking scenarios, and they struggle to synthesize animated frames with high-quality facial details and diverse domain styles.

### 2.2 Diffusion-based Portrait Animation

Diffusion models (DMs) [Ho et al. 2020; Song et al. 2020] achieves superior performance in various generative tasks including image generation [Rombach et al. 2021; Ruiz et al. 2023; Zhao et al. 2019] and editing [Brooks et al. 2023; Cao et al. 2023; Hertz et al. 2022], video generation [He et al. 2022b; Ma et al. 2024a,b; Singer et al. 2022] and editing [Liu et al. 2023c; Ma et al. 2023; Qi et al. 2023; Zhang et al. 2023b]. Recently, latent diffusion models further improved the performance by operating the diffusion step in latent space. Mainstream portrait animation approaches leverage the power of Stable Diffusion (SD) [Rombach et al. 2021] and incorporate temporal information into generation process, such as *AnimateDiff* [Guo et al. 2023], *MagicVideo* [Zhou et al. 2022], *VideoCrafter* [Chen et al. 2023] and *ModelScope* [Wang et al. 2023b]. Additionally, to preserve appearance context in the original image, many works [Chang et al. 2024; Hu et al. 2023; Xu et al. 2024b; Zhu et al. 2024] inject the reference image into the self-attention blocks in the LDM UNets, facilitating image editing and video generation. Even though these models can produce high-quality videos, they mainly rely on textual prompts for semantic guidance, which can be ambiguous and may not precisely describe users' intentions. To address such a problem, many control signals such as structure [gen 2023; King et al. 2024], pose [Chang et al. 2024; Ma et al. 2024a], and Canny edge [Zhang et al. 2023b] are applied for controllable video generation. Notely, several concurrent works [Chang et al. 2024; Hu et al. 2023; Xu et al. 2024b; Zhu et al. 2024] achieve state-of-the-art full-body by seamlessly integrating the appearance and motion control into temporal attentions. However, these methods all focus on full-body animation and ignore the specific details of the face. In contrast, we innovate the diffusion-based framework, focusing on driving various style portraits with detailed facial expressions (e.g., eyes, skins).

## 3 Method

The pipeline of our method is shown in Fig. 2. Given an input video clip, we randomly select a frame  $I_0$  as the reference portrait image. Then, we extract the motion sequences  $\{L_1, L_2, L_3, \dots, L_N\}$  (expression-aware landmarks) from the input video. The purpose of our method is to transfer the expression of the landmark sequences to  $I_0$ . Even for reference portraits of uncommon styles (i.e., cartoon,

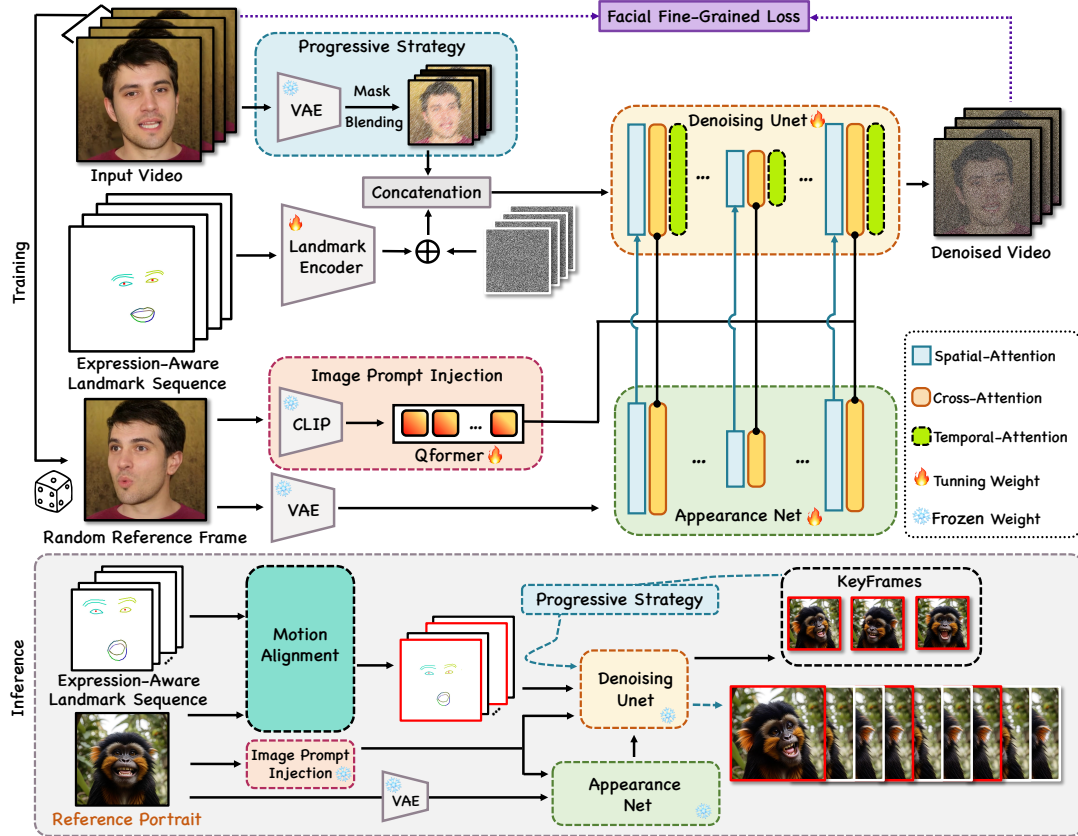


Fig. 2. **The overview of Follow-Your-Emoji.** We extract the features of our expression-aware landmark sequence with a landmark encoder and fuse these features with multi-frame noise first, then we utilize the progressive strategy to mask the frame of the input latent sequence randomly. Finally, we concatenate this latent sequence with the fused multi-frame noise and feed it to the Denoising UNet to conduct the denoising process for video generation. The appearance net and image prompt injection module help our model preserve the identity of the reference portrait, and the temporal attention maintains the temporal consistency. During training, the facial fine-grained loss guides the Unet to pay more attention to the facial and expression generation. During inference, we align the target landmark with the reference portrait with the motion alignment module. Then, we first generate the keyframes and utilize the progressive strategy to predict long videos.

sculpture, and animal), we hope our method can still predict good results.

We follow the recent diffusion-based portrait animation methods in our framework and utilize both the appearance net and temporal attention. For the control motions injection, we add the features of our expression-aware landmarks to UNet directly. These features are extracted with a landmark encoder. Moreover, similar to previous works [Liu et al. 2023b; Wang et al. 2024a,b; Ye et al. 2023], we encode the reference image  $I_0$  to image token using pre-trained CLIP image encoder, then the 4-layers Qformer [Zhang et al. 2023d] is employed to fuse all image token. In the next, we first discuss the motion representation and present our expression-aware landmark in Sec. 3.1. Then, we introduce the facial fine-grained loss in Sec. 3.2. Finally, for long-term animation, we describe the progressive strategy in Sec. 3.3.

### 3.1 Expression-Aware Landmark

Motion representation of facial expressions is essential for portrait animation. Accurate and precise motion representation enables

conveying the nuances of human emotion and expression, thereby enhancing the overall realism and impact of the animated portrait. Recent diffusion-based methods always directly utilize the portrait image sequences providing the driving motion [Xie et al. 2024] or the 2D landmarks as the motion representation for training. However, during the inference process, 2D landmarks cannot ensure alignment between the target expression and the reference portrait. This misalignment will lead to inaccurate generated expressions and potential leakage of the identity information. Directly using the portrait image providing the driving motion can solve this problem, but it is necessary to ensure that the person in the motion sequence is different from the reference portrait during the training process, which requires another portrait animation method for identity conversion. This conversion process will damage the accuracy of the expressions, and the portrait animation method can not transfer the identity of the uncommon portrait (i.e., turning a dog into a human).

To address the above problems, we introduce the expression-aware landmark, a new motion representation for portrait animation. Specifically, we utilize MediaPipe to extract the 3D keypoints of the




Fig. 3. **The detail of our facial fine-grained loss.** We extract the facial mask and expression mask with our landmark first. Then, we calculate the denoising loss  $\mathcal{L}_{FFG}$  in these masked regions.

portrait from the motion video. We then project these keypoints to obtain the 2D landmark. During the projection process, we discard the facial contour while retaining only the facial features. We find this operation can help the model focus on subtle motion generation and avoid the inaccuracy of facial contour with large expression changing, as shown in Fig. 7. Moreover, to capture the motion of the portrait’s irises, we calculate the related position of the irises in the eye sockets of 3D keypoints and maintain such a relationship after projection. In the end, since our expression-aware landmark is built on the 3D keypoints, we can align the target landmark sequence to the reference portrait in the canonical space of MediaPipe naturally, and we denote this process as motion alignment in the inference step as shown in Fig. 2.

### 3.2 Facial Fine-Grained Loss

For the portrait animation task, we hope the diffusion model focuses on expression generation and identity preservation. However, the diffusion model’s original training objective  $\mathcal{L}_{LDM}$  is to learn the content of all regions of the target image, which has no specific constraints for learning the facial content during the training process. Therefore, we propose the facial fine-grained (FFG) loss to modify the  $\mathcal{L}_{LDM}$  and make the model pay more attention to the content of facial and expression regions.

As shown in Fig. 3, we need to get two types of masks to capture the expression and facial regions to calculate the FFG loss. For the expression mask  $\mathcal{M}_e$ , we dilate each point of our expression-aware landmark and set these dilation regions as the expression mask. For the facial mask  $\mathcal{M}_f$ , we project the MediaPipe 3D facial contour’s keypoints and connect these projected points to get the facial masks. Finally, these two masks split the FFG loss into expression and facial aspects, respectively. Formally, the loss function can be written as below:

$$\mathcal{L}_{FFG} = \mathbb{E} \left[ \left\| \mathcal{M}_e \cdot (z - \hat{z}) + \mathcal{M}_f \cdot (z - \hat{z}) \right\|^2 \right] \quad (1)$$

where  $\hat{z}$  is the prediction latent embedding obtained by decoding the  $\epsilon_\theta$ . With our FFG loss, our method demonstrates better performance in both identity preservation and expression generation, as shown in Fig. 6. Finally, our total loss can be written as:

$$\mathcal{L} = \mathcal{L}_{LDM} + \mathcal{L}_{FFG} \quad (2)$$

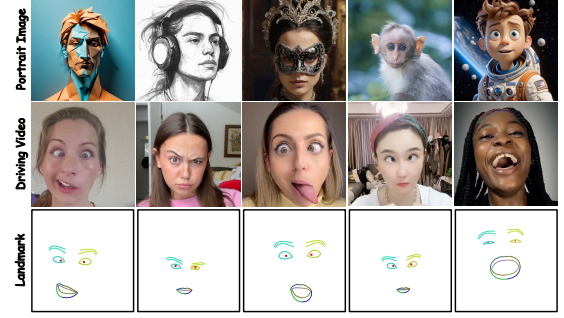


Fig. 4. **Examples of the EmojiBench with high expression diversity, exaggeration, and various visual styles in portrait images.**

### 3.3 Progressive Strategy for Long-Term Animation

With the advancement of technology and increasing user demands, long-term animation has become increasingly important in practical applications. Despite training on video clips, previous approaches [Chang et al. 2024; Hu et al. 2023; Xu et al. 2024b; Zhu et al. 2024] have also attempted to generate long videos during testing. They always synthesize several overlapping video clips and merge them using Gaussian smoothing. However, we observe that this trick leads to the degradation of temporal consistency.

To alleviate the above issues, a progressive strategy is proposed to generate long-term animation from coarse to fine. Intuitively, to generate the long-term animation in the inference step, we hope to generate keyframes first and then use these keyframes to generate the long-term animation with interpolation operation. To simulate this process, apart from the first and last latent frames, we mask the other input video latent frames first. Then, we concatenated this masked video latent with original UNet inputs to do the denoising process. With this strategy, we can set the first and last frames as keyframes in the inference step and help our model generate long-term animation. Meanwhile, we also mask each latent frame of the input video with a probability of 0.5, which helps our model generate the keyframes’s content in the first inference step since we need to cover all latent frames in this inference step. During the training process, we switch between these two masking strategies with a probability of 0.5.


## 4 Experiment

### 4.1 Implementation Details

We train our model on HDTF[Zhang et al. 2021], VFHQ[Xie et al. 2022], and our collected dataset jointly, which includes monocular camera recordings of 18 expressions and 20-minute real-human video from 115 subjects in indoor scenes. The outdoor scenes data is from VFHQ. The training stage consists of two stages, in the first training stage, individual video frames are sampled, resized, and center-cropped to a resolution of  $512 \times 512$ . We finetune it for 30k steps with a batch size of 32. In the second training stage, we train the temporal layer for 10,000 steps with 16-frame video sequences and a batch size of 32. The input images are from Civitai [civ 2023; Duc 2023; Wai 2023]. More details can be found in supplementary.

Table 1. **Quantitative comparisons with SOTA baselines.** We evaluate our framework both self and cross reenactments on  $256 \times 256$  test images.

Method	Self Reenactment				Cross Reenactment			User Study		
	L1 ↓	SSIM ↑	LPIPS ↓	FVD ↓	ID Similarity ↑	Image Quality ↑	Landmark Accuracy ↑	Expression ↓	Identity ↓	Overall ↓
Face Vid2vid [Wang et al. 2021]	0.042	0.801	0.244	228.9	0.627	39.568	6.12	6.46	7.81	5.92
DaGAN [Hong et al. 2022]	0.063	0.717	0.297	264.7	0.306	38.182	1.58	7.17	6.82	7.33
TPS [Zhao and Zhang 2022]	0.037	0.824	0.208	190.5	0.522	36.938	17.42	4.97	4.52	4.87
MCNet [Hong and Xu 2023]	0.031	0.837	0.193	208.1	0.412	33.981	18.63	5.52	5.31	5.92
FADM [Zeng et al. 2023]	0.045	0.701	0.268	187.4	0.679	33.496	3.95	3.16	3.78	3.86
MagicDance [Chang et al. 2024]	0.043	0.752	0.158	133.6	0.681	58.943	35.62	2.58	2.62	2.75
AniPortrait [Wei et al. 2024]	0.030	0.834	0.132	100.3	0.693	60.287	38.72	1.47	1.62	1.83
Ours	<b>0.029</b>	<b>0.849</b>	<b>0.136</b>	<b>96.8</b>	<b>0.702</b>	<b>66.287</b>	<b>39.16</b>	<b>1.26</b>	<b>1.45</b>	<b>1.72</b>

Fig. 5. **The qualitative comparisons with existing methods.** Given a reference portrait image and expression-aware landmarks, our approach demonstrates superior performance in capturing detailed facial expressions and maintaining the original identity of the characters compared to previous methods. More results are available in the supplementary material. The input images are from Civitai [civ 2023].Table 2. **Quantitative results of ablation study.** All metrics are evaluated on  $256 \times 256$  test images. ↑ indicates higher is better. ↓ indicates lower is better.

Method	Self Reenactment				Cross Reenactment		
	L1 ↓	SSIM ↑	LPIPS ↓	FVD ↓	ID Similarity ↑	Image Quality ↑	Landmark Accuracy ↓
FFG Loss (w/o Expression Mask)	0.037	0.702	0.159	147.4	0.576	53.792	26.87
FFG Loss (w/o Identity Mask)	0.036	0.721	0.157	149.3	0.548	50.992	34.21
w/o Progressive Strategy	0.035	0.718	0.141	138.8	0.632	52.108	35.32
2D Landmarks	0.039	0.715	0.166	144.1	0.521	50.829	17.45
w Facial Contour points	0.034	0.784	0.153	128.5	0.627	59.781	38.71
w/o Pupil points	0.035	0.762	0.147	103.2	0.648	61.436	33.58
Ours	<b>0.029</b>	<b>0.849</b>	<b>0.136</b>	<b>96.8</b>	<b>0.702</b>	<b>66.287</b>	<b>39.16</b>

## 4.2 EmojiBench

We introduce EmojiBench, a new benchmark to evaluate the model’s ability to animate freestyle portraits. Specifically, we collect 410 portraits from different domains, including cartoon style, real-human style, and even animals. These portrait cases are generated from 20 different personalized text-to-image models. We also provide 20 animal portraits, whose landmarks are able to be detected by Mediapipe [Lugaresi et al. 2019]. The EmojiBench contains 45 videos of driving human heads collected from the internet. Each video is approximately 5 seconds long with 150 frames. The expressions

of EmojiBench include a diverse range of head motion and facial expressions (e.g., frowning, crossed eyes, and pouting). Note that the subjects and expressions in training dataset and those in EmojiBench are different. The cases in EmojiBench do not appear in the training dataset. Such a benchmark with various styles would be beneficial for the development of the community.

## 4.3 Comparison with baselines

**4.3.1 Qualitative results.** We compare our approach with prior portrait animation methods visually, including state-of-the-art GAN-based methods Face Vid2vid [Wang et al. 2021], TPS [Zhao and

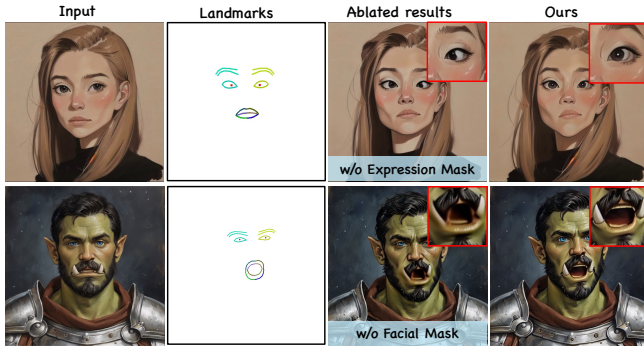


Fig. 6. **The effectiveness of facial fine-grained loss.** We analyze the performance of expression and facial aspects of FFG loss, respectively.

Zhang 2022], DaGAN [Hong et al. 2022], MCNet [Hong and Xu 2023] and recent diffusion-based including FADM [Zeng et al. 2023], MagicDance [Chang et al. 2024] and AniPortrait [Wei et al. 2024]. We exclude MegaPortraits [Drobyshev et al. 2022] and X-portrait [Xie et al. 2024] from our comparisons as no public release exists. We fine-tune all the baselines on our collected dataset. The results are shown in Fig. 5. We find the GAN-based method easily suffers from obvious artifacts, especially when changing the pose of the head with a large angle (i.e., see the generation result of the first character). Moreover, they can not rebuild the subtle expression for the reference portraits of uncommon style well (i.e., the movement of pupils in the second character). The diffusion-based methods MagicDance [Chang et al. 2024] and FADM [Zeng et al. 2023] perform better in expression transfer, but they still can not preserve the identity of reference portraits during animation. In contrast, our approach exhibits superior ability in handling large pose changing, subtle expressing generation, and identity preservation for uncommon style portraits. See more results in Fig. 8 and Fig. 9.

**4.3.2 Quantitative results.** We compare our method with state-of-the-art portrait animation on our EmojiBench quantitatively and the results are shown in Tab. 1. Due to the limited resolution of most previous works, all measurements are performed in 64 frames at a resolution of  $256 \times 256$ . All evaluation metrics used are as follows: (a) **Self reenactment:** For quantitative assessment of image-level quality, we report the four metrics, L1 error, SSIM [Wang et al. 2004], LPIPS [Zhang et al. 2018], and FVD [Unterthiner et al. 2018]. For each video in EmojiBench, the first frame is employed as the reference image to generate the facial expression sequences. We leverage subsequent frames to serve as both the driving image and the ground truth. (b) **Cross reenactment** means that reference images and landmark sequences are from different subjects. We evaluate cross reenactment on four metrics: identity similarity, image quality, expression landmark accuracy, and user study, respectively. (1) *Identity similarity:* the ArcFace score [Deng et al. 2019] is applied to measure identity preservation. We calculate cosine similarity between source and generated images. (2) *Image quality assessment:* We follow [Xie et al. 2024] to utilize the HyperIQA [Zhang et al. 2023c] for image quality assessment. (3) *Landmark accuracy:* To evaluate the pose accuracy of the generated video, we regard the input facial landmark

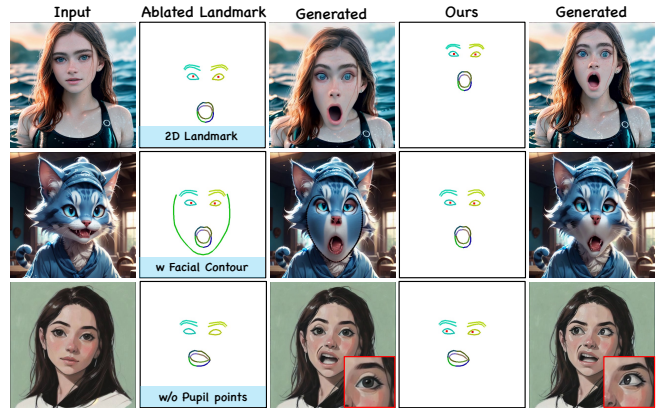


Fig. 7. **The effectiveness of expression-aware landmark.** We compare the results when different landmarks is used to guide the portrait animation.

sequences as ground truth and detect the 2D landmark of generated portraits, computing the average accuracy with driving landmark sequences. (c) **User study:** we perform the user study on cross reenactment with three aspects. (1) *Expression:* Evaluating the quality of generated expression. (2) *Identity:* Measuring the identity similarity between the generated frame images and input reference portrait image. (3) *Overall:* Evaluating the overall quality of the generated videos. See more details in supplementary. As reported in Tab. 1, our approach achieves the best performance in seven metrics of self/cross reenactment and user study metrics.

#### 4.4 Ablation Study

In the subsequent section, we will analyze the effectiveness of expression-aware landmark and facial fine-grained loss. As for progressive strategy for long-term animation, we provide more discussion in supplementary materials. Note that we perform all ablations in same training setting.

*Effectiveness of expression-aware landmark.* To prove the effectiveness of our expression-aware landmark, we change our motion representation to the 2D landmark, expression-aware landmark with the facial contour, and expression-aware landmark without pupil points to generate the video, respectively. The visual results are shown in Fig. 7. 2D landmark has a challenge in handling the alignment of the facial bounding box between target landmarks and reference portrait images, as presented in the 1st row. The expression-aware landmark with the facial contour fails to maintain the identity of portrait images in non-human styles. This is because the current open-source landmark detector makes it hard to predict the facial contour of any style portrait. Finally, we also show the result produced with expression-aware landmarks without pupil points. Due to the lack of motion signals of pupil points, it is difficult to generate lively expressions with pupil motion. In contrast, our full model demonstrates better performance. The corresponding numerical evaluation is shown in Tab. 2.

*Effectiveness of facial fine-grained loss.* To analyze the performance of FFG loss, we discarded the expression and facial aspects of FFG loss separately to do the experiment. Without facial aspects of FFG loss, we find our method reduces the ability to protect identity information and detail appearance of the input portrait (i.e., teeth disappeared in the second row of Fig. 6). Meanwhile, when we abandon the expression aspects of FFG loss, our method can not capture the subtle expression changing well (i.e., inaccurate pupil movement in the first row of Fig. 6). The corresponding numerical evaluation is shown in Tab. 2.

## 5 Discussion

*Conclusion.* We introduce Follow-Your-Emoji, a novel diffusion-based framework for freestyle portrait animation. Incorporating with the expression-aware landmark, our method shows high performance in subtle and exaggerated facial expression generation. Meanwhile, we propose a facial fine-grained loss to constrain the diffusion model focus on expression generation and identity preservation. To train our model, we introduce a new expression training dataset with 18 exaggerated expressions and 20-minute real-human videos from 115 subjects. Then, we introduce the progressive strategy for stable long-term animation. Finally, to address the lack of benchmark in portrait animation, we build the EmojiBench, a comprehensive benchmark to evaluate our method, the impressive performance of our model on generalized reference portraits and driving motions serves as validation of its effectiveness.

## Acknowledgments

This project was supported by the National Key R&D Program of China under grant number 2022ZD0161501.



## References

2023. civitai. <https://civitai.com/models/443821/cyberrealistic-pony>.
2023. duchaitenpony-real. <https://civitai.com/models/477851/duchaiten-pony-real>.
2023. Gen-2. <https://runwayml.com/ai-magic-tools/gen-2/>.
2023. wairealmix. <https://civitai.com/models/393905/wai-realmix>.
- Hadar Averbuch-Elor, Daniel Cohen-Or, Johannes Kopf, and Michael F Cohen. 2017. Bringing portraits to life. *ACM transactions on graphics (TOG)* 36, 6 (2017), 1–13.
- Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. 2023. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 22563–22575.
- Tim Brooks, Aleksander Holynski, and Alexei A Efros. 2023. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 18392–18402.
- Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiaohu Qie, and Yinqiang Zheng. 2023. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 22560–22570.
- Di Chang, Yichun Shi, Quankai Gao, Jessica Fu, Hongyi Xu, Guoxian Song, Qing Yan, Yizhe Zhu, Xiao Yang, and Mohammad Soleymani. 2024. MagicPose: Realistic Human Poses and Facial Expressions Retargeting with Identity-aware Diffusion. [arXiv:2311.12052](https://arxiv.org/abs/2311.12052) [cs.CV]
- Haoxin Chen, Menghan Xia, Yingqing He, Yong Zhang, Xiaodong Cun, Shaoshu Yang, Jinbo Xing, Yaofang Liu, Qifeng Chen, Xintao Wang, Chao Weng, and Ying Shan. 2023. VideoCrafter1: Open Diffusion Models for High-Quality Video Generation. [arXiv:2310.19512](https://arxiv.org/abs/2310.19512) [cs.CV]
- Haoxin Chen, Yong Zhang, Xiaodong Cun, Menghan Xia, Xintao Wang, Chao Weng, and Ying Shan. 2024. VideoCrafter2: Overcoming Data Limitations for High-Quality Video Diffusion Models. [arXiv:2401.09047](https://arxiv.org/abs/2401.09047) [cs.CV]
- Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. 2019. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 4690–4699.
- Yu Deng, Jiaolong Yang, Dong Chen, Fang Wen, and Xin Tong. 2020. Disentangled and controllable face image generation via 3d imitative-contrastive learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 5154–5163.
- Nikita Drobyshev, Jenya Chelishev, Taras Khakhulin, Aleksei Ivakhnenko, Victor Lempitsky, and Egor Zakharov. 2022. Megaportraits: One-shot megapixel neural head avatars. In *Proceedings of the 30th ACM International Conference on Multimedia*. 2663–2671.
- Yao Feng, Haiwen Feng, Michael J Black, and Timo Bolkart. 2021. Learning an animatable detailed 3D face model from in-the-wild images. *ACM Transactions on Graphics (ToG)* 40, 4 (2021), 1–13.
- Ohad Fried, Ayush Tewari, Michael Zollhöfer, Adam Finkelstein, Eli Shechtman, Dan B Goldman, Kyle Genova, Zeyu Jin, Christian Theobalt, and Maneesh Agrawala. 2019. Text-based editing of talking-head video. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–14.
- Yuan Gong, Yong Zhang, Xiaodong Cun, Fei Yin, Yanbo Fan, Xuan Wang, Baoyuan Wu, and Yujiu Yang. 2023. Toontalker: Cross-domain face reenactment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7690–7700.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (2020), 139–144.
- Yuwei Guo, Ceyuan Yang, Anyi Rao, Yaohui Wang, Yu Qiao, Dahua Lin, and Bo Dai. 2023. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. [arXiv preprint arXiv:2307.04725](https://arxiv.org/abs/2307.04725) (2023).
- Tianyu He, Junliang Guo, Runyi Yu, Yuchi Wang, Jialiang Zhu, Kaikai An, Leyi Li, Xu Tan, Chunyu Wang, Han Hu, et al. 2023. GAIA: Zero-shot Talking Avatar Generation. [arXiv preprint arXiv:2311.15230](https://arxiv.org/abs/2311.15230) (2023).
- Yingqing He, Tianyu Yang, Yong Zhang, Ying Shan, and Qifeng Chen. 2022a. Latent Video Diffusion Models for High-Fidelity Long Video Generation. (2022). [arXiv:2211.13221](https://arxiv.org/abs/2211.13221) [cs.CV]
- Yingqing He, Tianyu Yang, Yong Zhang, Ying Shan, and Qifeng Chen. 2022b. Latent video diffusion models for high-fidelity long video generation. [arXiv preprint arXiv:2211.13221](https://arxiv.org/abs/2211.13221) (2022).
- Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. 2022. Prompt-to-prompt image editing with cross attention control. [arXiv preprint arXiv:2208.01626](https://arxiv.org/abs/2208.01626) (2022).
- Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. 2022a. Imagen video: High definition video generation with diffusion models. [arXiv preprint arXiv:2210.02303](https://arxiv.org/abs/2210.02303) (2022).
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems* 33 (2020), 6840–6851.
- Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. 2022b. Video diffusion models. *Advances in Neural Information Processing Systems* 35 (2022), 8633–8646.
- Fa-Ting Hong and Dan Xu. 2023. Implicit Identity Representation Conditioned Memory Compensation Network for Talking Head video Generation. In *ICCV*.
- Fa-Ting Hong, Longhao Zhang, Li Shen, and Dan Xu. 2022. Depth-aware generative adversarial network for talking head video generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 3397–3406.
- Li Hu, Xin Gao, Peng Zhang, Ke Sun, Bang Zhang, and Liefeng Bo. 2023. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. [arXiv preprint arXiv:2311.17117](https://arxiv.org/abs/2311.17117) (2023).
- Taras Khakhulin, Vanessa Sklyarova, Victor Lempitsky, and Egor Zakharov. 2022. Realistic one-shot mesh-based head avatars. In *European Conference on Computer Vision*. Springer, 345–362.
- Hyeonwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Niessner, Patrick Pérez, Christian Richardt, Michael Zollhöfer, and Christian Theobalt. 2018. Deep video portraits. *ACM transactions on graphics (TOG)* 37, 4 (2018), 1–14.
- Gongye Liu, Menghan Xia, Yong Zhang, Haoxin Chen, Jinbo Xing, Xintao Wang, Yujiu Yang, and Ying Shan. 2023b. StyleCrafter: Enhancing Stylized Text-to-Video Generation with Style Adapter. [arXiv preprint arXiv:2312.00330](https://arxiv.org/abs/2312.00330) (2023).
- Hongyu Liu, Xingtong Han, Chengbin Jin, Lihui Qian, Huawei Wei, Zhe Lin, Faqiang Wang, Haoye Dong, Yibing Song, Jia Xu, et al. 2023a. Human motionformer: Transferring human motions with vision transformers. [arXiv preprint arXiv:2302.11306](https://arxiv.org/abs/2302.11306) (2023).
- Shaoteng Liu, Yuechen Zhang, Wenbo Li, Zhe Lin, and Jiaya Jia. 2023c. Video-p2p: Video editing with cross-attention control. [arXiv preprint arXiv:2303.04761](https://arxiv.org/abs/2303.04761) (2023).
- Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Guang Yong, Juhyun Lee, et al. 2019. Mediapipe: A framework for building perception pipelines. [arXiv preprint arXiv:1906.08172](https://arxiv.org/abs/1906.08172) (2019).
- Yue Ma, Xiaodong Cun, Yingqing He, Chenyang Qi, Xintao Wang, Ying Shan, Xiu Li, and Qifeng Chen. 2023. MagicStick: Controllable Video Editing via Control Handle Transformations. [arXiv preprint arXiv:2312.03047](https://arxiv.org/abs/2312.03047) (2023).
- Yue Ma, Yingqing He, Xiaodong Cun, Xintao Wang, Siran Chen, Xiu Li, and Qifeng Chen. 2024a. Follow your pose: Pose-guided text-to-video generation using pose-free videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 4117–4125.
- Yue Ma, Yingqing He, Hongfa Wang, Andong Wang, Chenyang Qi, Chengfei Cai, Xiu Li, Zhifeng Li, Heung-Yeung Shum, Wei Liu, et al. 2024b. Follow-Your-Click: Open-domain Regional Image Animation via Short Prompts. [arXiv preprint arXiv:2403.08268](https://arxiv.org/abs/2403.08268) (2024).
- Chenyang Qi, Xiaodong Cun, Yong Zhang, Chenyang Lei, Xintao Wang, Ying Shan, and Qifeng Chen. 2023. Fatezero: Fusing attentions for zero-shot text-based video editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 15932–15942.
- Linzi Qu, Jiaxiang Shang, Xiaoguang Han, and Hongbo Fu. 2023. ReenactArtFace: Artistic Face Image Reenactment. *IEEE Transactions on Visualization and Computer Graphics* (2023).
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. PMLR, 8748–8763.
- Aditya Ramesh, Pratul Dharwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical text-conditional image generation with clip latents. [arXiv preprint arXiv:2204.06125](https://arxiv.org/abs/2204.06125) 1, 2 (2022), 3.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2021. High-Resolution Image Synthesis with Latent Diffusion Models. [arXiv:2112.10752](https://arxiv.org/abs/2112.10752) [cs.CV]
- Nataníel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 2023. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 22500–22510.
- 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. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems* 35 (2022), 36479–36494.
- Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. 2019. First order motion model for image animation. *Advances in neural information processing systems* 32 (2019).
- Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gefni, et al. 2022. Make-a-video: Text-to-video generation without text-video data. [arXiv preprint arXiv:2209.14792](https://arxiv.org/abs/2209.14792) (2022).
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*. PMLR, 2256–2265.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. 2020. Denoising diffusion implicit models. [arXiv preprint arXiv:2010.02502](https://arxiv.org/abs/2010.02502) (2020).

- Jingxiang Sun, Xuan Wang, Lizhen Wang, Xiaoyu Li, Yong Zhang, Hongwen Zhang, and Yebin Liu. 2023. Next3d: Generative neural texture rasterization for 3d-aware head avatars. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 20991–21002.
- Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. 2016. Face2face: Real-time face capture and reenactment of rgb videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2387–2395.
- Linrui Tian, Qi Wang, Bang Zhang, and Liefeng Bo. 2024. EMO: Emote Portrait Alive-Generating Expressive Portrait Videos with Audio2Video Diffusion Model under Weak Conditions. *arXiv preprint arXiv:2402.17485* (2024).
- Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. 2018. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717* (2018).
- Haofan Wang, Qixun Wang, Xu Bai, Zekui Qin, and Anthony Chen. 2024a. InstantStyle: Free Lunch towards Style-Preserving in Text-to-Image Generation. *arXiv preprint arXiv:2404.02733* (2024).
- Haofan Wang, Peng Xing, Renyuan Huang, Hao Ai, Qixun Wang, and Xu Bai. 2024b. InstantStyle-Plus: Style Transfer with Content-Preserving in Text-to-Image Generation. *arXiv preprint arXiv:2407.00788* (2024).
- Jiuniu Wang, Hangjie Yuan, Dayou Chen, Yingya Zhang, Xiang Wang, and Shiwei Zhang. 2023. Modelscope text-to-video technical report. *arXiv preprint arXiv:2308.06571* (2023).
- Tan Wang, Linjie Li, Kevin Lin, Yuanhao Zhai, Chung-Ching Lin, Zhengyuan Yang, Hanwang Zhang, Zicheng Liu, and Lijuan Wang. 2023a. Disco: Disentangled control for realistic human dance generation. *arXiv preprint arXiv:2307.00040* (2023).
- Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. 2021. One-Shot Free-View Neural Talking-Head Synthesis for Video Conferencing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612.
- Huawei Wei, Zejun Yang, and Zhisheng Wang. 2024. Aniportrait: Audio-driven synthesis of photorealistic portrait animation. *arXiv preprint arXiv:2403.17694* (2024).
- Olivia Wiles, A Koepke, and Andrew Zisserman. 2018. X2face: A network for controlling face generation using images, audio, and pose codes. In *Proceedings of the European conference on computer vision (ECCV)*. 670–686.
- Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. 2023. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7623–7633.
- Liangbin Xie, Xintao Wang, Honglun Zhang, Chao Dong, and Ying Shan. 2022. 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*. 657–666.
- You Xie, Hongyi Xu, Guoxian Song, Chao Wang, Yichun Shi, and Linjie Luo. 2024. X-Portrait: Expressive Portrait Animation with Hierarchical Motion Attention. *arXiv preprint arXiv:2403.15931* (2024).
- Jinbo Xing, Menghan Xia, Yuxin Liu, Yuechen Zhang, Y He, H Liu, H Chen, X Cun, X Wang, Y Shan, et al. 2024. Make-Your-Video: Customized Video Generation Using Textual and Structural Guidance. *IEEE Transactions on Visualization and Computer Graphics* (2024).
- Hongyi Xu, Guoxian Song, Zihang Jiang, Jianfeng Zhang, Yichun Shi, Jing Liu, Wanchun Ma, Jiashi Feng, and Linjie Luo. 2023. Omniavatar: Geometry-guided controllable 3d head synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12814–12824.
- Sicheng Xu, Guojun Chen, Yu-Xiao Guo, Jialong Yang, Chong Li, Zhenyu Zang, Yizhong Zhang, Xin Tong, and Baining Guo. 2024a. VASA-1: Lifelike Audio-Driven Talking Faces Generated in Real Time. *arXiv preprint arXiv:2404.10667* (2024).
- Zhongcong Xu, Jianfeng Zhang, Jun Hao Liew, Hanshu Yan, Jia-Wei Liu, Chenxu Zhang, Jiashi Feng, and Mike Zheng Shou. 2024b. MagicAnimate: Temporally Consistent Human Image Animation using Diffusion Model.
- Zhendong Yang, Ailing Zeng, Chun Yuan, and Yu Li. 2023. Effective whole-body pose estimation with two-stages distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4210–4220.
- Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. 2023. IP-Adapter: Text Compatible Image Prompt Adapter for Text-to-Image Diffusion Models. (2023).
- Bohan Zeng, Xuhui Liu, Sicheng Gao, Boyu Liu, Hong Li, Jianzhuang Liu, and Baochang Zhang. 2023. Face animation with an attribute-guided diffusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 628–637.
- Qiming Zhang, Jing Zhang, Yufei Xu, and Dacheng Tao. 2023d. Vision Transformer with Quadrangle Attention. *arXiv preprint arXiv:2303.15105* (2023).
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 586–595.
- Wenxuan Zhang, Xiaodong Cun, Xuan Wang, Yong Zhang, Xi Shen, Yu Guo, Ying Shan, and Fei Wang. 2023a. 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*. 8652–8661.
- Weixia Zhang, Guangtao Zhai, Ying Wei, Xiaokang Yang, and Kede Ma. 2023c. Blind Image Quality Assessment via Vision-Language Correspondence: A Multitask Learning Perspective. In *IEEE Conference on Computer Vision and Pattern Recognition*. 14071–14081.
- Yabo Zhang, Yuxiang Wei, Dongsheng Jiang, Xiaopeng Zhang, Wangmeng Zuo, and Qi Tian. 2023b. Controlvideo: Training-free controllable text-to-video generation. *arXiv preprint arXiv:2305.13077* (2023).
- Zhimeng Zhang, Lincheng Li, Yu Ding, and Changjie Fan. 2021. 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*. 3661–3670.
- Bo Zhao, Lili Meng, Weidong Yin, and Leonid Sigal. 2019. Image generation from layout. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8584–8593.
- Jian Zhao and Hui Zhang. 2022. Thin-plate spline motion model for image animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 3657–3666.
- Daquan Zhou, Weimin Wang, Hanshu Yan, Weiwei Lv, Yizhe Zhu, and Jiashi Feng. 2022. Magicvideo: Efficient video generation with latent diffusion models. *arXiv preprint arXiv:2211.11018* (2022).
- Shenhao Zhu, Junming Leo Chen, Zuoqiuo Dai, Yinghui Xu, Xun Cao, Yao Yao, Hao Zhu, and Siyu Zhu. 2024. Champ: Controllable and Consistent Human Image Animation with 3D Parametric Guidance. *arXiv:2403.14781* [cs.CV]

Driving



Animation Results




Fig. 8. More portrait animation results.

Driving

Animation Results

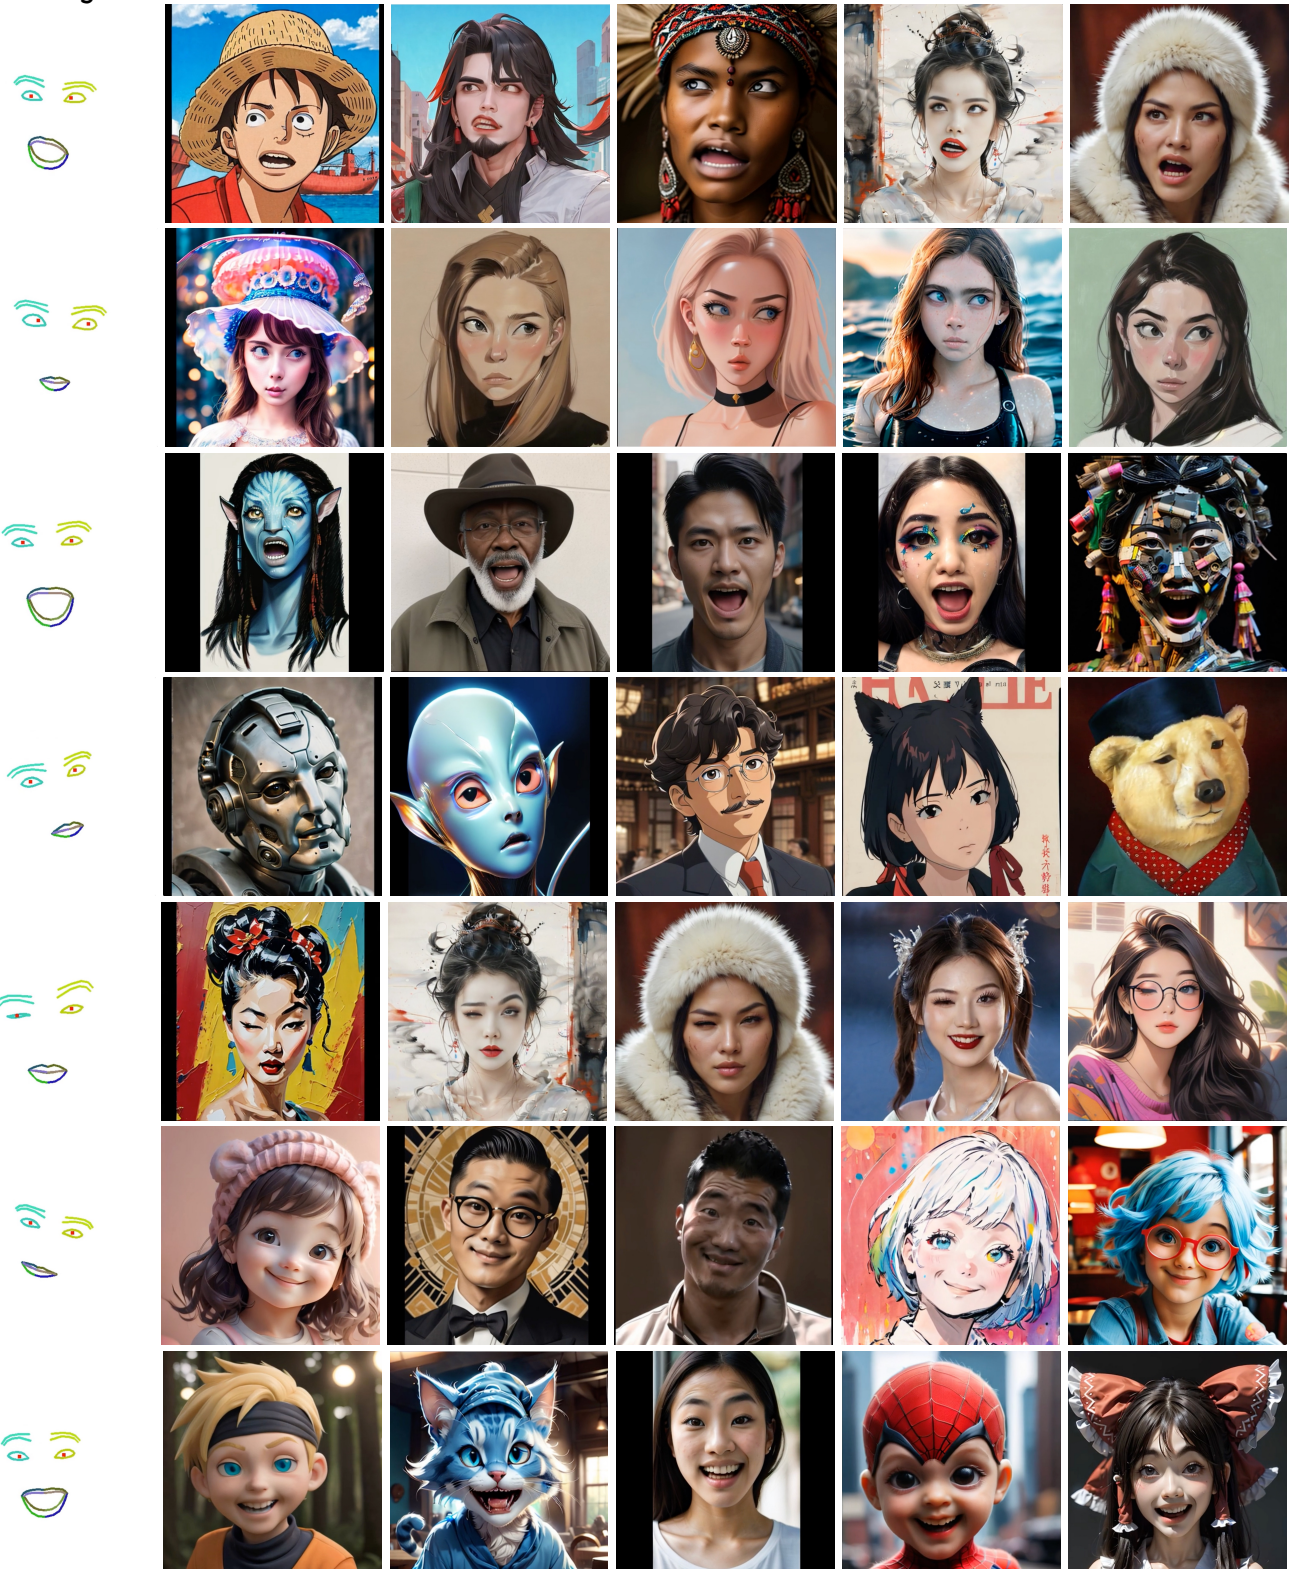


Fig. 9. More portrait animation results.