

---

# SkyReels-A1: Expressive Portrait Animation in Video Diffusion Transformers

---

Di Qiu, Zhengcong Fei, Rui Wang, Jialin Bai, Changqian Yu  
Mingyuan Fan, Guibin Chen, Xiang Wen

Skywork AI

Project page: [skyreels-a1.github.io](https://skyreels-a1.github.io)




Figure 1: **SkyReels-A1 can generate an animated portrait from a reference image, driven by video motions while ensuring identity preservation.** Notably, our method ensures accurate transfer of facial expressions and body movements, allowing for realistic, high-quality portrait animations at anybody proportion that integrate naturally into different scenes.

## Abstract

We present SkyReels-A1<sup>1</sup>, a simple yet effective framework built upon video diffusion Transformers (DiT) to facilitate portrait image animation. Existing methodologies still encounter issues, including identity distortion, background instability, and unrealistic facial dynamics, particularly in head-only animation scenarios. Besides, extending to accommodate diverse body proportions usually leads to visual inconsistencies or unnatural articulations. To address these challenges, SkyReels-A1 capitalizes on the strong generative capabilities of video DiT, enhancing facial motion transfer precision, identity retention, and temporal coherence. The system incorporates an expression-aware conditioning module that enables seamless video synthesis driven by expression-guided landmark inputs. Integrating the facial image-text alignment module strengthens the fusion of facial

---

<sup>1</sup>Core contributors: {qiuidihk, feizhengcong}@gmail.com.

attributes with motion trajectories, reinforcing identity preservation. Additionally, SkyReels-A1 incorporates a multi-stage training paradigm to incrementally refine the correlation between expressions and motion while ensuring stable identity reproduction. Extensive empirical evaluations highlight the model’s ability to produce visually coherent and compositionally diverse results, making it highly applicable to domains such as virtual avatars, remote communication, and digital media generation. The code and model weights of SkyReels-A1 are publicly available at <https://github.com/SkyworkAI/SkyReels-A1>.

## 1 Introduction

Generation of dynamic facial movements and expressions from a single static portrait, known as portrait animation, has emerged as a prominent research area due to its fascinating and diverse applications, including virtual avatars, digital communication, and narrative-driven media. Recent progress, especially with diffusion-based techniques[27, 49, 50], has significantly advanced the domain. Earlier approaches [44, 45, 48, 56, 62, 9, 47] predominantly relied on generative adversarial networks (GANs) [19] to animate head movements, employing facial landmark tracking or 3D morphable models [61] as motion guidance mechanisms. More recent developments[29, 59, 4] have integrated diffusion models to improve temporal stability and realism, with large pre-trained video generative models [57, 38, 6, 13, 12]. It demonstrates superior capabilities in capturing intricate facial dynamics[2, 60, 16], and enhances both pose-driven facial synthesis and audio-driven talking head generation. Moreover, commercial applications such as Runway’s Act-One have achieved notable progress in producing expressive and dynamic animated portraits, achieving lifelike facial expressions while adapting performances to characters with varying proportions.

Although it marks a substantial leap in the application of generative models for creative media production[57, 6, 38, 30, 55, 54], producing high-quality animations that maintain natural expressiveness and identity consistency across diverse body compositions remains an unresolved challenge [15], particularly when animation extends beyond head movements, i.e., diverse body proportions. Existing methods [38, 22] are constrained by issues such as identity drift, unstable background generation, and overly rigid facial expressions, which become more pronounced in full-body animation scenarios [20]. These challenges largely arise from the difficulty of seamlessly integrating identity-preserving features with motion dynamics, the absence of fine-grained conditioning for nuanced expressions, and limitations in training methodologies for comprehensive motion adaptation.

This study introduces SkyReels-A1, an advanced framework built upon the video diffusion Transformers (DiT) architecture, aimed at producing high-fidelity and expressive portrait animations. Our approach is distinguished by three primary innovations: (i) We harness the representational capacity of DiT to jointly model intricate facial details and full-body dynamics within a unified latent space, effectively addressing challenges such as identity drift and background inconsistencies. (ii) We incorporate an expression-aware conditioning module, which fuses continuous landmark sequences with facial image-text alignment features, facilitating precise manipulation of subtle expressions, *e.g.*, eyebrow elevation, and lip curvature while ensuring identity coherence. (iii) A multi-stage training paradigm is implemented, which iteratively enhances pose accuracy, identity stability, and motion realism in sequence, leading to more natural and lifelike animations.

By unifying these components, SkyReels-A1 achieves superior performance in generating portrait animations that maintain identity integrity, exhibit natural expressiveness, and seamlessly accommodate variations in body proportions. Extensive empirical evaluations indicate that our framework surpasses existing approaches in both quantitative assessments and user studies, particularly in handling complex anatomical structures and intricate micro-expressions. Additionally, the modular design of SkyReels-A1 facilitates straightforward integration into downstream applications, including real-time video editing systems and personalized avatar platforms. This work represents a significant advancement in portrait animation by reconciling identity preservation with motion fidelity, thereby offering a robust framework for next-generation digital human technologies.

In summary, our contributions are listed as:

- We present SkyReels-A1, a novel framework for portrait animation that employs a DiT architecture to enhance fidelity in motion transfer precision, identity retention, and temporal

coherence. The framework incorporates a dynamic conditioning module informed by expression landmarks and a cross-modal alignment mechanism that bridges visual-textual semantic spaces.

- A phased training methodology is formulated to incrementally refine motion-expression correlation and subject-specific feature invariance. Initial stages establish coarse motion-texture associations, while subsequent phases employ adversarial regularization to synchronize fine-grained facial dynamics with latent identity embeddings.
- A comprehensive series of experiments has been conducted to evaluate the performance, which indicates that SkyReels-A1, yields highly effective outcomes, exhibiting seamless adaptability to a wide range of compositional variations. This versatility renders our method particularly well-suited for various applications, including but not limited to the creation of virtual avatars, enhancement of video conferencing experiences, and the production of digital content. Finally, to support further research and industry adoption, both the code and demonstration materials are publicly available.

## 2 Related Works

### 2.1 GAN-based Portrait Animation

Early approaches to portrait animation have utilized GANs for learning motion dynamics in a self-supervised manner. The seminal works [46, 47] typically involve a two-stage process: warping and rendering. These methods use sparse neural keypoints to predict head motion, which is then applied to warp the encoded identity features of the source. The warped features are subsequently passed through a generative decoder to produce animated frames, incorporating in-painted facial details and background elements. In subsequent research, several approaches have sought to improve the warping estimation, with notable contributions including the use of 3D neural keypoints [56], depth information [28], and thin-plate splines [48]. Additionally, ReenactArtFace [40] employed a 3D morphable model to control expressions and poses, while ToonTalker [18] applied a transformer-based architecture to enhance the warping process, particularly in the context of cross-domain datasets. Concurrently, efforts to refine the rendering process have resulted in several advancements. For instance, MegaPortraits [9] elevated the resolution of the generated portraits to megapixels by utilizing high-resolution image datasets. FADM [62] proposed a coarse-to-fine animation framework, incorporating a diffusion process to enhance facial details from the initial outputs generated by FOMM [47] and Face Vid2Vid [56]. Beyond traditional video reenactment, significant strides have been made in leveraging alternative driving signals, such as 3D facial priors [8, 32, 51, 17, 58] and audio cues [23, 52, 11]. Despite these advancements, existing methods primarily rely on explicit feature warping, with an emphasis on facial animation in talking scenarios.

### 2.2 Diffusion-based Portrait Animation

Recent advancements in diffusion models [27, 49, 50] have led to exceptional performance across a variety of generative tasks, including image [43], video [25, 2, 1], and multi-view renderings [36, 35, 21]. Latent diffusion models (LDMs) [42, 14] have further refined this approach by reducing computational costs, performing the diffusion process within a lower-dimensional latent space. In the context of portrait animation, pre-trained diffusion models [43, 42] have been widely adopted for image-to-video tasks. Notably, several studies [3, 34] have demonstrated the effectiveness of integrating reference image features into the self-attention blocks of LDM UNets, enabling improved image editing and video generation while preserving the original appearance context. In addition, ControlNet [64] extends the LDM architecture by incorporating structural control signals, such as landmarks, segmentations, and dense poses, to guide controllable image generation. This has inspired several concurrent works [29, 59, 4, 39, 10] that achieve state-of-the-art full-body portrait animation by seamlessly integrating appearance and motion controls along with temporal attention mechanisms [25, 24] within pre-trained UNet models. Despite these advances, the motion control signals employed in these methods—such as skeletons with or without facial landmarks [29, 4] and dense poses [59]—may not fully capture the original motion dynamics. Furthermore, recent approaches [38, 6] that focus on facial expression and motion control have introduced more expressive landmarks for enhanced facial representation. However, these methods face challenges when applied to diverse body proportions, as they struggle to maintain consistent facial dynamics when transitioning

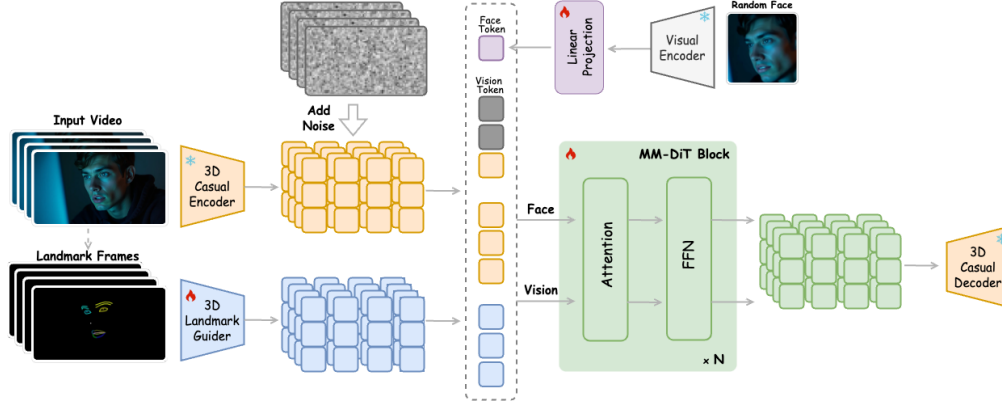


Figure 2: **Overview of SkyReels-A1 framework.** Given an input video sequence and a reference portrait image, we extract facial expression-aware landmarks from the video, which serve as motion descriptors for transferring expressions onto the portrait. Utilizing a conditional video generation framework based on DiT, our approach directly integrates these facial expression-aware landmarks into the input latent space. In alignment with prior research, we employ a pose guidance mechanism constructed within a VAE architecture. This component encodes facial expression-aware landmarks as conditional input for the DiT framework, thereby enabling the model to capture essential low-dimensional visual attributes while preserving the semantic integrity of facial features.

to half-body or full-body compositions. In this work, we investigate the use of an advanced video diffusion transformers architecture to enhance portrait animation capabilities.

### 3 Methodology

#### 3.1 Preliminary

Diffusion models generate structured data by progressively reversing a noise corruption process, transforming random noise into coherent samples. At each time step  $t \in \{0, \dots, T\}$ , the model estimates the latent state  $x_t$  conditioned on  $x_{t+1}$ , as described by:

$$p_{\theta}(x_t|x_{t+1}) = \mathcal{N}(x_t; \tilde{\mu}_t, \tilde{\beta}_t I), \quad (1)$$

where  $\theta$  denotes the model parameters,  $\tilde{\mu}_t$  is the predicted mean, and  $\tilde{\beta}_t$  is the variance schedule. The training objective typically utilizes a mean squared error (MSE) loss on the noise prediction  $\hat{\epsilon}_{\theta}(x_t, t, c_{\text{txt}})$  as:

$$\mathcal{L}_{\text{noise}} = w \cdot \mathbb{E}_{t, c_{\text{txt}}, \epsilon \sim \mathcal{N}(0,1)} \left[ \|\epsilon - \epsilon_{\theta}(x_t, t, c_{\text{txt}})\|^2 \right], \quad (2)$$

where  $c_{\text{txt}}$  denotes the text condition. Transformer-based text-to-video diffusion models have shown significant promise in generating video content. Our proposed model, SkyReels-A1, leverages an advanced transformer-based latent diffusion framework, incorporating a 3D Variational Autoencoder (VAE) to map videos from the pixel domain into a latent space. Each fundamental transformer block within the model consists of a 3D spatial-temporal self-attention mechanism and a feed-forward network (FFN). The text embedding  $c_{\text{txt}}$  is obtained by the T5 encoder. Recent studies on controllable generation usually extend this framework by incorporating additional control signals, such as image condition  $c_{\text{img}}$ . It is achieved through a feature extractor  $\tau_{\text{img}}$  that processes a reference image  $r$ :  $c_{\text{img}} = \tau_{\text{img}}(r)$  and the noise prediction function in 2 becomes  $\epsilon_{\theta}(x_t, t, c_{\text{txt}}, c_{\text{img}})$ .

#### 3.2 SkyReels-A1 Framework

**Overview.** The overall architecture of SkyReels-A1 is depicted in Figure 2. Given an input video clip and a reference portrait image  $I_{\text{ref}}$ , we extract facial expression-aware landmarks from the video as motion sequences to transfer the captured expressions onto the portrait. Based on a DiT-based conditional video generation framework, our model adopts a direct approach by incorporating

the features of our facial expression-aware landmarks into the input latent. Similar to previous works, We employed a pose guider based on a VAE architecture to encode facial expression-aware landmarks as conditional input for DiT, ensuring that the model captures accurate low-dimensional visual information while preserving facial semantic information. In between, the Facial Image-Text Alignment module enhances identity consistency by mapping facial features into the text feature space, enabling precise control over facial expressions and motion while preserving the character’s identity. Specifically, we first use a face extractor  $\mathcal{E}_{face}$  to locate the facial region  $I_{face}$ :

$$I_{face} = \mathcal{E}_{face}(I_{ref}). \tag{3}$$

To better extract the overall semantic information of the facial image  $I_{face}$ , facial features are obtained using a high-performance vision encoder  $\mathcal{E}_{vision}$ , i.e., SigLip[63]:

$$F_{id} = \mathcal{E}_{vision}(I_{face}). \tag{4}$$

Through the aforementioned strategies and architectural design, our method facilitates the generation of high-quality, expression-driven video content with improved realism and coherence. In the following, we will describe the main module in SkyReels-A1 in detail.

**Expression-Aware Landmark.** Inspired by [38, 6], we introduce a similar facial landmark-based motion condition input for portrait animation. Accurate motion representation of facial expressions constitutes a fundamental requirement for expressive portrait animation, as it captures the nuanced spectrum of human emotions and micro-expressions that significantly enhance the realism and emotional resonance of animated avatars. Current diffusion-based approaches predominantly employ 2D facial landmarks as motion representations during training. However, both paradigms exhibit critical limitations: The reliance on 2D landmarks during inference frequently results in misalignment between target expressions and reference portraits, manifesting as expression mismatches and identity leakage artifacts. While alternative methods leverage third-party tools like MediaPipe to extract 3D keypoints from video sequences, such representations often lack the precision required to capture fine-grained expression details and complex facial dynamics, particularly in non-frontal orientations and extreme expressions.

To overcome the aforementioned challenges, we propose 3D Facial Expressions[41], a framework that integrates a neural rendering module to enhance the accuracy and realism of reconstructed facial expressions. Unlike conventional approaches that rely on differentiable rendering, our method replaces this component with a neural rendering mechanism, facilitating a more effective learning process and improved generalization across diverse facial expressions. This architecture enables the extraction of high-precision 3D keypoints, capturing intricate motion details and complex facial dynamics with greater fidelity. By leveraging this refined motion representation, our approach significantly enhances the realism of portrait animations while ensuring improved expression accuracy, identity consistency, and adaptability across various application scenarios.

**3D Landmark Guider.** To ensure spatio-temporal coherence between the driving signals and the input video latent representations, we propose the Spatio-temporal Alignment Landmark Guide Module. Specifically, it utilizes a 3D causal encoder as the core component for landmark guidance. Fine-tuned carefully, it can capture the driving signals’ motion representation more effectively, ensuring precise and coherent alignment between the motion signals and the input video latents.

By leveraging the 3D causal encoder, the Spatio-temporal Alignment Landmark Guide module directly projects the driving signals into a shared latent space with the video latents. This shared representation bridges the gap between the motion signals and the generated video frames, resulting in synchronized spatial and temporal dynamics. The fine-tuning process further refines the encoder’s ability to capture intricate motion patterns, thereby enhancing the realism and fidelity of motion transfer. This approach not only ensures precise motion alignment but also guarantees the preservation of both identity and motion consistency in the generated video, enabling high-quality, temporally stable animations.

**Facial Image-Text Alignment.** In existing portrait animation methods, maintaining consistent identity while altering facial expressions remains a challenging problem that requires further investigation. Specifically, earlier approaches have focused on enhancing identity consistency using identity-preserving adapters through cross-attention mechanisms. However, we found that this method is difficult to train and introduces a significant number of additional parameters. Inspired by the

architecture of CogVideoX[60], we aim to improve facial identity consistency during the expression generation process by concatenating the embeddings of both facial images and video at the input stage. This approach not only enhances identity consistency but also allows for the seamless transfer of the pre-trained base model’s capabilities. To achieve this, we introduce a lightweight learnable mapping module through a multi-layer perceptron  $\mathcal{P}$ , which maps facial features into the text feature space as:

$$F'_{id} = \mathcal{P}(F_{id}), \tag{5}$$

where  $F_{id}$  is the identity embedding extracted by vision encoder  $\mathcal{E}_{vision}$  [63], a pre-trained image-text model. Since  $F_{id}$  captures detailed facial features while visual encompasses broader facial information and is less sensitive to external factors such as lighting and occlusion, the fusion of both types of information enhances the accuracy of facial features in the generated results. The Face Image-Text Alignment Module facilitates better information integration while maintaining model efficiency.

## 4 Training

### 4.1 Progressive Training Strategy

The training pipeline of SkyReels-A1 is structured into three distinct stages: Motion-Driven Training, Identity-Preserving Training, and Multi-module Joint Fine-Tuning. Each stage plays a pivotal role in progressively refining the model’s ability to generate high-fidelity, temporally consistent character animations.

- **Motion-Driven Training.** At this stage, we aim to incorporate motion conditions into the video generation process. Motion-driven video inputs are processed using a 3D landmark guidance module, which is initialized from a pre-trained 3D causal encoder but remains unchanged during this phase. To effectively integrate landmark-based motion information while maintaining the integrity of the image-to-video (IT2V) architecture, the extracted landmark latent representation is concatenated with the noise input. Model adaptation to landmark-specific variations is achieved by training only the convolutional layers within the PatchEmbedding module, ensuring that the system refines motion representation while preserving the core IT2V functionalities.
- **Identity-Preserving Training.** In the subsequent stage, the focus shifts to reinforcing identity preservation in the animated portrait. Although textual conditions are not explicitly required for portrait animation, the text-processing branch is retained to capitalize on the pre-trained model’s capabilities. Facial features are encoded using a CLIP image encoder[63] to generate a facial representation, which is subsequently projected into the text feature space through a trainable linear mapping. Notably, only this projection layer undergoes optimization, while all other components remain unchanged. This design choice ensures that the animated portrait maintains identity consistency across varying expressions and motion sequences.
- **Multi-module Joint Fine-Tuning.** This stage ensures the model’s ability to generate precise, high-quality animations by jointly optimizing the 3D landmark guider, DiT block, and the linear projection layer. It further improves the model’s generalization across diverse portrait types and motion scenarios, enabling it to handle a wide range of facial expressions and dynamic movements with greater fidelity.

### 4.2 Face-Aware Loss

The fluency of the generated video depends largely on the spatial coherence and realism of dynamic regions, such as the face. To enhance this, we apply a face-aware loss that directs the model to prioritize high-motion areas.

Here we use RAFT [53] as  $\Theta$  for efficient and accurate optical flow estimation. Then the mean optical flow value  $\tau_i$  can be calculated by simply averaging  $f_{i,(x,y)}$ . Afterward, we take  $\tau_i$  as the threshold to produce binary mask  $\mathcal{M}_{i,(x,y)}$ . Specifically, when the magnitude of the optical flow exceeds  $\tau_i$ , set the corresponding position in  $\mathcal{M}_{i,(x,y)}$  to 1; otherwise set it to 0. Consequently, the mean foreground



Figure 3: **Qualitative portrait animation results from SkyReels-A1.** Given a static portrait image as input, our model can vividly animate it, ensuring seamless stitching and offering precise control over eyes and lip movements.

optical flow value  $f_{i,fg}$  can be easily obtained by:

$$f_{i,fg} = \frac{1}{S} \sum_{x=1}^H \sum_{y=1}^W f_{i,fg}(x, y) = \frac{1}{S} \sum_{x=1}^H \sum_{y=1}^W \mathcal{M}_{i,(x,y)} * f_{i,(x,y)}, \quad (6)$$

where  $f_{i,fg}(x, y)$  is the foreground optical flow at each pixel  $(x, y)$ .  $S$  denotes the number of the foreground pixels. We then normalize the foreground optical flow  $f_{i,fg}(x, y)$  defined in Equation (6) and calculate the optical flow mask  $\mathcal{M}_{i,norm}$  as following:

$$\mathcal{M}_{i,norm} = \text{clip}(f_{i,fg}(x, y)/255 + 0.5, 1.0, 1.5), \quad (7)$$

where function  $\text{clip}(\cdot, a, b)$  restricts the values into  $[a, b]$ . The high-motion areas will be assigned a greater value than the low-motion regions. Then the face-aware loss  $L_{face}$  across all  $N$  frames can be compactly defined as:

$$L_{face} = \frac{1}{NH'W'} \sum_{i=1}^N \sum_{x=1}^{H'} \sum_{y=1}^{W'} \mathcal{M}_{i,norm} \cdot [\epsilon_i(x, y) - \hat{\epsilon}_i(x, y)]^2, \quad (8)$$

where  $\epsilon_i(x, y)$  and  $\hat{\epsilon}_i(x, y)$  denote the target and predicted noise at location  $(x, y)$ , respectively.  $H'$  and  $W'$  correspond to the resolution of latent features.

## 5 Experiment

We first give an overview of the implementation details, including data sources, data filtering processes, baselines, and benchmarks utilized in our experiments. Next, we present comparative experimental results against selected baselines, to validate the effectiveness of the proposed modules.

### 5.1 Experimental Settings

**Implementation Details.** We train our model based on the advanced video diffusion transformer CogVideoX-5B[60] using a combination of our collected dataset and publicly available datasets. During the multi-stage training phases, we train for 2K steps in the first stage, followed by another 2K steps in the second stage, and conclude with 1K steps in the final stage, using a batch size of 512. The learning rate is set to  $1e-5$  for the first two stages and reduced to  $1e-6$  in the final stage, utilizing AdamW optimizers. Our experiments are conducted on 32 NVIDIA A800 GPUs. During inference, we utilize DDIM sampler and set the scale of classifier-free guidance to 3 in our experiment. The static images used as reference are generated by Flux and sourced from the Pexels.

Method	ID Similarity $\uparrow$	Image Quality $\downarrow$	Expression Dis. $\downarrow$	Pose Dis. $\downarrow$
Follow-Your-Emoji [38]	0.5771 / 0.5983	100.2192	0.0417	0.8502
LivePortrait [22]	0.7011 / 0.7305	83.3168	0.0396	0.8372
Act-One	0.7219 / 0.7470	68.5953	0.0329	0.8204
<b>SkyReels-A1</b>	0.7196 / 0.7314	<b>59.6884</b>	0.0363	0.8245

Table 1: **Quantitative comparisons of portrait animation.** SkyReels-A1 exhibits a substantial enhancement in image quality, outperforming all existing approaches by a notable margin, with the exception of the proprietary commercial model Act-One.



Figure 4: **Qualitative comparisons results.** Our SkyReels-A1 model better transfers lip movements and eye gazes from another person, while maintaining the identity of the source portrait.

**Dataset Sources.** The training video clips featured are sourced from the NeRSemble[33] dataset, HDTF[65], DFEW[31], RAVDESS[37], Panda70M[5] and our collected dataset that we collected approximately 1W character video clips from the internet.

**Dataset Filtering.** During the data pre-processing stage, we implement a series of meticulous filtering steps to ensure the quality and suitability of the video-text dataset. The workflow consists of three phases: single-character extraction, motion filtering, and post-processing. First, we select single-character videos to clean the video content using existing tools, addressing issues such as camera capture artifacts and background noise. Next, we extract head pose information and mouth landmarks using facial key points detected by MediaPipe. By calculating head angles and mouth variations, we filter samples with significant facial expressions and head movements. Finally, based on the detected facial positions from the previous steps, We crop or pad the videos to a fixed resolution of  $480 \times 720$  to meet the model’s input requirements. A random frame is then selected from each video, and the face is encoded into an embedding using clip encoder[63], providing essential facial feature information for the model.

**Baselines.** To comprehensively assess SkyReels-A1’s performance across diverse scenarios, we compare it with portrait animation baselines, including the open-source solution LivePortrait[22], Follow-Your-Emoji[38], and the closed-source commercial product Runway Act One<sup>2</sup>.

<sup>2</sup><https://runwayml.com/research/introducing-act-one>






Figure 5: More generated results from Skyreels-A1 in diverse body proportions.

**Evaluation Metrics.** To measure the generalization quality and motion accuracy of portrait animation results, we employ three metrics to evaluate identity similarity, image quality, and expression and head pose accuracy, respectively. Specifically, the FaceSim-Arc and FaceSim-Cur score [7], calculating the cosine similarity between source and generated images, is utilized to assess identity preservation. We employ a pre-trained network with FID [26] for image quality assessment. To

evaluate the motion accuracy, we compare the L1 difference between the extracted facial blendshapes and head poses of the driving and generated frames using FaceAnalysis<sup>3</sup> and OpenPose<sup>4</sup>.

## 5.2 Comparison with Baselines

**Quantitative Results.** We conducted experiments to investigate cross-identity motion transfer, wherein a reference portrait was randomly selected from a set of 100 in-the-wild images, while the driving sequence was obtained from our test dataset. Table 1 presents the quantitative evaluation results. We can see that our proposed model demonstrates superior performance compared to both diffusion-based and non-diffusion-based approaches in terms of generation fidelity and motion precision. Notably, by incorporating video diffusion transformers as a prior, SkyReels-A1 achieves a significant improvement in image quality, surpassing existing methods by a considerable margin except the closed-source commercial model Act-One.

**Qualitative Results.** Figure 4 presents qualitative comparisons of portrait animation, complementing the results obtained from automated evaluation metrics. The first two examples highlight the model’s robustness in transferring motion accurately, even under significant pose variations in the driving or source portraits. In the third and fourth cases, the model effectively captures and transfers subtle facial expressions, such as lip movements and eye gazes, while preserving the visual consistency of the original portrait. Furthermore, the final case demonstrates that the integration of stitching techniques enhances the model’s stability, allowing it to animate full-body images seamlessly, even when the reference portrait contains a relatively small facial region.

## 5.3 User Study

To further validate the superiority of the proposed SkyReels-A1 model over prior baselines in motion accuracy and expressiveness, a user study was conducted. Specifically, 20 participants from diverse geographic regions were recruited to evaluate synthesized videos. Each participant responded to a series of comparative questions assessing two key aspects: motion accuracy and human likeness. The evaluation was conducted with explicit knowledge of the model names, and participants were instructed to select the video that most accurately replicated the expressions and movements of the driving sequence. Among the 100 collected responses, 63% indicated a preference for SkyReels-A1, affirming its enhanced capability in preserving both facial expressions and pose fidelity compared to existing baselines.

## 6 Conclusion

In this study, we present SkyReels-A1, an innovative portrait animation framework based on video diffusion transformers. By integrating motion and identity representations, our approach achieves high-fidelity generation of both subtle and exaggerated facial expressions. Through extensive automatic and user-based evaluations, we have demonstrated the robustness and adaptability of our model across various customized scenarios. We anticipate that these promising results will contribute to advancements in portrait animation applications.

**Limitations.** Despite its strong performance, the current model exhibits limitations when handling extreme pose variations. Addressing these challenges remains a key focus for future research.

---

<sup>3</sup><https://github.com/deepinsight/insightface>

<sup>4</sup><https://github.com/CMU-Perceptual-Computing-Lab/openpose>

## References

- [1] Yuxiang Bao, Di Qiu, Guoliang Kang, Baochang Zhang, Bo Jin, Kaiye Wang, and Pengfei Yan. Latentwarp: Consistent diffusion latents for zero-shot video-to-video translation. *arXiv preprint arXiv:2311.00353*, 2023.
- [2] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, Varun Jampani, and Robin Rombach. Stable video diffusion: Scaling latent video diffusion models to large datasets, 2023.
- [3] Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiaohu Qie, and Yinqiang Zheng. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22560–22570, 2023.
- [4] Di Chang, Yichun Shi, Quankai Gao, Jessica Fu, Hongyi Xu, Guoxian Song, Qing Yan, Yizhe Zhu, Xiao Yang, and Mohammad Soleymani. Magicpose: Realistic human poses and facial expressions retargeting with identity-aware diffusion, 2024.
- [5] Tsai-Shien Chen, Aliaksandr Siarohin, Willi Menapace, Ekaterina Deyneka, Hsiang-wei Chao, Byung Eun Jeon, Yuwei Fang, Hsin-Ying Lee, Jian Ren, Ming-Hsuan Yang, et al. Panda-70m: Captioning 70m videos with multiple cross-modality teachers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13320–13331, 2024.
- [6] 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.
- [7] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, pages 4690–4699, 2019.
- [8] Yu Deng, Jiaolong Yang, Dong Chen, Fang Wen, and Xin Tong. Disentangled and controllable face image generation via 3d imitative-contrastive learning. In *IEEE Computer Vision and Pattern Recognition*, 2020.
- [9] 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*, 2022.
- [10] Haopeng Fang, Di Qiu, Binjie Mao, Pengfei Yan, and He Tang. Motioncharacter: Identity-preserving and motion controllable human video generation. *arXiv preprint arXiv:2411.18281*, 2024.
- [11] Zhengcong Fei, Mingyuan Fan, Changqian Yu, and Junshi Huang. Flux that plays music. *arXiv preprint arXiv:2409.00587*, 2024.
- [12] Zhengcong Fei, Mingyuan Fan, Changqian Yu, and Junshi Huang. Scalable diffusion models with state space backbone. *arXiv preprint arXiv:2402.05608*, 2024.
- [13] Zhengcong Fei, Mingyuan Fan, Changqian Yu, Debang Li, and Junshi Huang. Scaling diffusion transformers to 16 billion parameters. *arXiv preprint arXiv:2407.11633*, 2024.
- [14] Zhengcong Fei, Mingyuan Fan, Changqian Yu, Debang Li, Youqiang Zhang, and Junshi Huang. Dimba: Transformer-mamba diffusion models. *arXiv preprint arXiv:2406.01159*, 2024.
- [15] Zhengcong Fei, Debang Li, Di Qiu, Changqian Yu, and Mingyuan Fan. Ingredients: Blending custom photos with video diffusion transformers. *arXiv preprint arXiv:2501.01790*, 2025.
- [16] Zhengcong Fei, Di Qiu, Changqian Yu, Debang Li, Mingyuan Fan, and Xiang Wen. Video diffusion transformers are in-context learners. *arXiv preprint arXiv:2412.10783*, 2024.
- [17] Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. Learning an animatable detailed 3D face model from in-the-wild images. volume 40, 2021.
- [18] Yuan Gong, Yong Zhang, Xiaodong Cun, Fei Yin, Yanbo Fan, Xuan Wang, Baoyuan Wu, and Yujiu Yang. Toontalker: Cross-domain face reenactment, 2023.
- [19] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.

- [20] Shreyank N Gowda, Dheeraj Pandey, and Shashank Narayana Gowda. From pixels to portraits: A comprehensive survey of talking head generation techniques and applications. *arXiv preprint arXiv:2308.16041*, 2023.
- [21] Yuming Gu, You Xie, Hongyi Xu, Guoxian Song, Yichun Shi, Di Chang, Jing Yang, and Linjie Luo. Diffportrait3d: Controllable diffusion for zero-shot portrait view synthesis, 2023.
- [22] 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.
- [23] Yudong Guo, Keyu Chen, Sen Liang, Yongjin Liu, Hujun Bao, and Juyong Zhang. Ad-nerf: Audio driven neural radiance fields for talking head synthesis. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [24] Yuwei Guo, Ceyuan Yang, Anyi Rao, Maneesh Agrawala, Dahua Lin, and Bo Dai. Sparsectrl: Adding sparse controls to text-to-video diffusion models. *arXiv preprint arXiv:2311.16933*, 2023.
- [25] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. *arXiv preprint arXiv:2307.04725*, 2023.
- [26] 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. In *NeurIPS*, 2017.
- [27] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [28] Fa-Ting Hong, Longhao Zhang, Li Shen, and Dan Xu. Depth-aware generative adversarial network for talking head video generation. 2022.
- [29] Li Hu, Xin Gao, Peng Zhang, Ke Sun, Bang Zhang, and Liefeng Bo. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. *arXiv preprint arXiv:2311.17117*, 2023.
- [30] 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.
- [31] Xingxun Jiang, Yuan Zong, Wenming Zheng, Chuangao Tang, Wanchuang Xia, Cheng Lu, and Jiateng Liu. Dfew: A large-scale database for recognizing dynamic facial expressions in the wild. In *Proceedings of the 28th ACM international conference on multimedia*, pages 2881–2889, 2020.
- [32] Taras Khakhulin, Vanessa Sklyarova, Victor Lempitsky, and Egor Zakharov. Realistic one-shot mesh-based head avatars. In *European Conference of Computer vision (ECCV)*, 2022.
- [33] Tobias Kirschstein, Shenhan Qian, Simon Giebenhain, Tim Walter, and Matthias Nießner. Nersemble: Multi-view radiance field reconstruction of human heads. *ACM Transactions on Graphics (TOG)*, 42(4):1–14, 2023.
- [34] Yukang Lin, Haonan Han, Chaoqun Gong, Zunnan Xu, Yachao Zhang, and Xiu Li. Consistent123: One image to highly consistent 3d asset using case-aware diffusion priors, 2023.
- [35] Minghua Liu, Chao Xu, Haiyan Jin, Linghao Chen, Mukund Varma T, Zexiang Xu, and Hao Su. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization, 2023.
- [36] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object, 2023.
- [37] Steven R Livingstone and Frank A Russo. The ryerson audio-visual database of emotional speech and song (ravdess): A dynamic, multimodal set of facial and vocal expressions in north american english. *PloS one*, 13(5):e0196391, 2018.
- [38] Yue Ma, Hongyu Liu, Hongfa Wang, Heng Pan, Yingqing He, Junkun Yuan, Ailing Zeng, Chengfei Cai, Heung-Yeung Shum, Wei Liu, et al. Follow-your-emoji: Fine-controllable and expressive freestyle portrait animation. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–12, 2024.

- [39] Di Qiu, Zheng Chen, Rui Wang, Mingyuan Fan, Changqian Yu, Junshi Huang, and Xiang Wen. Moviecharacter: A tuning-free framework for controllable character video synthesis. *arXiv preprint arXiv:2410.20974*, 2024.
- [40] Linzi Qu, Jiayang Shang, Xiaoguang Han, and Hongbo Fu. Reenactartface: Artistic face image reenactment. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–13, 2023.
- [41] George Retsinas, Panagiotis P Filntisis, Radek Danecek, Victoria F Abrevaya, Anastasios Roussos, Timo Bolkart, and Petros Maragos. 3d facial expressions through analysis-by-neural-synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2490–2501, 2024.
- [42] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021.
- [43] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022.
- [44] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. Animating arbitrary objects via deep motion transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2377–2386, 2019.
- [45] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. In *NeurIPS*, 2019.
- [46] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. Animating arbitrary objects via deep motion transfer. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [47] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. December 2019.
- [48] 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.
- [49] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- [50] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- [51] Jingxiang Sun, Xuan Wang, Lizhen Wang, Xiaoyu Li, Yong Zhang, Hongwen Zhang, and Yebin Liu. Next3d: Generative neural texture rasterization for 3d-aware head avatars. In *CVPR*, 2023.
- [52] Yasheng Sun, Hang Zhou, Kaisiyuan Wang, Qianyi Wu, Zhibin Hong, Jingtuo Liu, Errui Ding, Jingdong Wang, Ziwei Liu, and Koike Hideki. Masked lip-sync prediction by audio-visual contextual exploitation in transformers. In *SIGGRAPH Asia 2022 Conference Papers*, 2022.
- [53] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 402–419. Springer, 2020.
- [54] Linrui Tian, Qi Wang, Bang Zhang, and Liefeng Bo. Emo: Emote portrait alive generating expressive portrait videos with audio2video diffusion model under weak conditions. In *European Conference on Computer Vision*, pages 244–260. Springer, 2024.
- [55] Cong Wang, Kuan Tian, Jun Zhang, Yonghang Guan, Feng Luo, Fei Shen, Zhiwei Jiang, Qing Gu, Xiao Han, and Wei Yang. V-express: Conditional dropout for progressive training of portrait video generation. *arXiv preprint arXiv:2406.02511*, 2024.
- [56] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *CVPR*, 2021.
- [57] You Xie, Hongyi Xu, Guoxian Song, Chao Wang, Yichun Shi, and Linjie Luo. X-portrait: Expressive portrait animation with hierarchical motion attention. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024.

- [58] H. Xu, G. Song, Z. Jiang, J. Zhang, Y. Shi, J. Liu, W. Ma, J. Feng, and L. Luo. Omniaavatar: Geometry-guided controllable 3d head synthesis. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [59] Zhongcong Xu, Jianfeng Zhang, Jun Hao Liew, Hanshu Yan, Jia-Wei Liu, Chenxu Zhang, Jiashi Feng, and Mike Zheng Shou. Magicanimate: Temporally consistent human image animation using diffusion model. 2023.
- [60] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024.
- [61] Dan Ye and Chiou-Shann Fuh. 3d morphable face model for face animation. *International Journal of Image and Graphics*, 20(01):2050003, 2020.
- [62] Bohan Zeng, Xuhui Liu, Sicheng Gao, Boyu Liu, Hong Li, Jianzhuang Liu, and Baochang Zhang. Face animation with an attribute-guided diffusion model. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2023.
- [63] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. in 2023 ieee. In *CVF International Conference on Computer Vision (ICCV)*, pages 11941–11952, 2023.
- [64] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023.
- [65] 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.