FaceLift: Learning Generalizable Single Image 3D Face Reconstruction from Synthetic Heads

Weijie Lyu[1*] Yi Zhou[2] Ming-Hsuan Yang[1] Zhixin Shu[2]
[1]University of California, Merced [2]Adobe Research

Figure 1. *FaceLift* transforms a single facial image into a high-fidelity 3D Gaussian head representation. Trained exclusively on synthetic 3D data, our pipeline first generates sparse, identity-preserving multiview images of the input head using a diffusion model. These sparse generated views are then fed into a transformer-based 3D Gaussian splats reconstructor, producing complete and detailed 3D head representation that generalize remarkably well to real-world human images. Project page: https://weijielyu.github.io/FaceLift.

Abstract

We present FaceLift, a novel feed-forward approach for generalizable high-quality 360-degree 3D head reconstruction from a single image. Our pipeline first employs a multi-view latent diffusion model to generate consistent side and back views from a single facial input, which then feed into a transformer-based reconstructor that produces a comprehensive 3D Gaussian splats representation. Previous methods for monocular 3D face reconstruction often lack full view coverage or view consistency due to insufficient multi-view supervision. We address this by creating a high-quality synthetic head dataset that enables consistent supervision across viewpoints. To bridge the domain gap between synthetic training data and real-world images, we propose a simple yet effective technique that ensures the view generation process maintains fidelity to the input by learning to reconstruct the input image alongside the view generation. Despite being trained exclusively on synthetic data, our method demonstrates remarkable generalization to real-world images. Through extensive qualitative and quantitative evaluations, we show that FaceLift outperforms state-of-the-art 3D face reconstruction methods on identity preservation, detail recovery and rendering quality.

1. Introduction

3D face reconstruction has been a central focus in computer vision and graphics research for decades, driven by its crucial applications in immersive virtual and augmented realities, VFX and gaming, digital entertainment, and next-generation telepresence systems. However, achieving high quality reconstruction from a single image remains very challenging. On one hand, the monocular face reconstruction problem is highly ill-posed – a single 2D image can be produced by countless different 3D face shapes, creating fundamental ambiguity. On the other hand, the human visual system is highly attuned to facial details, making even subtle artifacts and imperfections noticeable to the eye.

Traditional 3D head synthesis approaches typically use parametric textured mesh models [32, 60] trained on 3D scan datasets. While these models enable basic head generation, the rendered images frequently lack fine-scale geometric detail, realistic textures, and convincing hair, limiting their perceptual realism and expressiveness. Recent breakthroughs in image generative models [19, 23] and novel view synthesis techniques [27, 40] have opened new possibilities for this research area. Leveraging these developments, recent works [1, 72] use neural 3D representations to learn effective 3D head representation from unstructured real face image datasets [25, 68]. While these datasets improve the realism and diversity of rendering results, they fail to provide multi-view supervision for modeling 3D consis-

*Work was done when Weijie Lyu was an intern at Adobe Research.

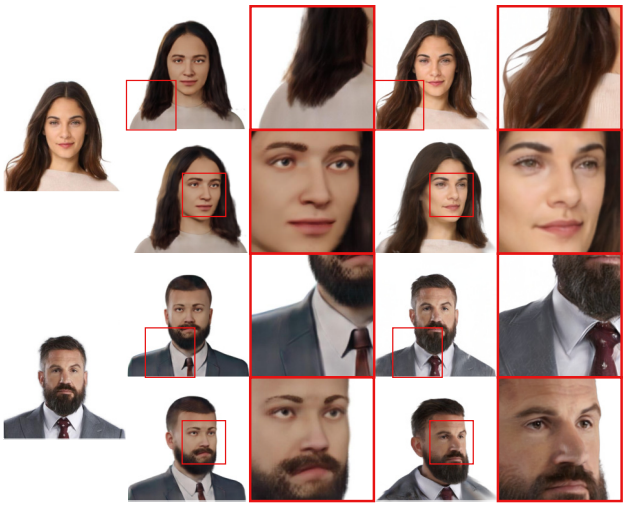

Figure 2. **Comparison with RodinHD.** RodinHD [73] trains tri-plane diffusion with synthetic data, resulting in apparent identity loss. In contrast, *FaceLift* achieves better identity preservation and generalizes effectively to real human portraits.

tency causing view inconsistency and artifacts on the back of the head. Since diverse multi-view real images are hard to acquire, RodinHD [73] leverages synthetic multi-view images to train generative models that directly output 3D neural representations of the head. However, training solely on synthetic data often results in significant perceptual identity loss in the generated outputs, as demonstrated in Fig. 2.

In this work, we present *FaceLift*, a pipeline for learning generalizable and high-fidelity single image to 3D face representation from synthetic head data. To achieve high quality reconstruction that preserves the input facial identity, we adopt a two-stage pipeline to first generate identity preserving multi-view images using a diffusion model [48], followed by a transformer-based feed-forward reconstructor to fuse the generated sparse views into a comprehensive 3D Gaussian representation. We train the model with synthetic images – multi-view renderings of 3D synthetic human portraits using Blender. We highlight two key techniques for generalizing to real-world images and preserving input facial identity. First, we emphasize the importance of reconstructing the input image alongside the view synthesis task in the conditional diffusion model training, which significantly improved generalization capability in testing. Second, we demonstrate that training the feed-forward reconstructor benefits from a two-stage training process: pre-training on general objects [12] to acquire a rich geometry and texture prior, followed by fine-tuning on synthetic human head data to capture head-specific geometry. With our two-stage approach, we focus on learning identity preservation in the image space during the first stage, achieving higher input fidelity compared to existing methods.

Comparing with prior art, we achieve three key advancements: (1) robust view consistency through multi-view attention and supervision, (2) improved generalization from training techniques and foundational model, ensuring accurate identity preservation, (3) high-quality facial texture and hair details via pixel-aligned Gaussian representation.

We extensively evaluate *FaceLift* quantitatively and qualitatively across diverse datasets. Using real multi-view studio captures [39] and an independent synthetic human dataset [8], our approach consistently surpasses previous state-of-the-art methods across all evaluation metrics. Through extensive testing on in-the-wild portrait images, we demonstrate that our method reconstructs complete 3D heads with fine-grained details, accurate identity preservation, and high visual fidelity. Comparisons and ablation studies confirm that multi-view consistent training data is crucial for high-fidelity face reconstruction. Our contributions are summarized as follows:

- We propose *FaceLift*, a framework that reconstructs a high-fidelity 3D head from a single image using view generation and feed-forward reconstructor.
- Despite being trained solely on synthetic human head data, our approach shows no domain gap on real-world images, highlighting both the effectiveness of synthetic data and our model's robust generalization capabilities.
- We construct two benchmarks on single-image to 3D full head reconstruction tasks using the publicly available datasets Cafca [8] and Ava-256 [39] to quantitatively evaluate models' performance on both reconstruction accuracy and identity preservation ability.
- Our comprehensive evaluation confirms that our approach achieves state-of-the-art performance in reconstruction accuracy and identity preservation.

2. Related Work

Face Reconstruction. 3D face reconstruction has been a long-standing challenge in computer vision, with substantial progress driven by diverse approaches. Vetter and Blanz [60] pioneer a method for synthesizing 3D faces by linearly blending multiple 3D templates, now widely known as blendshapes. This work establishes the foundation for 3D Morphable Models (3DMMs), which represent 3D face shapes and textures as principal components derived from scanned data. Subsequent research [5, 6, 32, 34, 47] extend this framework, enabling the generation of new 3D faces by manipulating blending coefficients. However, these methods produce mesh-based representations that lack fine details and are limited to modeling the front of the face, excluding hair and 360-degree synthesis. While 3DMM-based methods have been foundational, recent advances in deep learning, especially Generative Adversarial Networks (GANs) [19, 25, 26], have greatly improved 3D face synthesis quality. EG3D [72] uses a tri-plane NeRF representation with a pose-conditioned StyleGAN2 [26] framework. Follow-up works [3, 33] achieve single-image-to-3D gen-

eration through GAN inversion [11]. Despite their success, these methods can only synthesize near-frontal views. To overcome this, PanoHead [1] introduces a tri-grid neural volume representation, enabling full 360-degree head synthesis. Unfortunately, it does not provide a 3D head representation for consistent multi-view rendering. Recent efforts explore alternative representations for 3D face reconstruction from sparse input, such as a single image [17, 42, 61] or few-shot images [7]. However, these methods still require pre-instance optimization. Rodin [63] and its extension RodinHD [73] employ an image-conditioned diffusion model to generate a triplane representation of a human head for full-head novel view synthesis. Nevertheless, their triplane diffusion model is limited to synthetic data and struggles to achieve high-fidelity reconstructions from real-world images. For animatable 3D head avatars generations, Morphable Diffusion [10] generates multi-view consistent images from a single image using a morphable mesh, while HeadGAP [76] generates 3D animatable head avatars using few-shot input, leveraging 3D head priors learned from large-scale data. In contrast, our work focuses on leveraging synthetic training data to produce high-fidelity, detailed 3D Gaussian head models.

Synthetic Human Data. Capturing high-quality 3D data of real humans requires a controlled studio environment and costly photography equipment [39]. As an alternative, large-scale synthetic 3D head datasets have emerged as an effective and resource-efficient solution for tasks like human head reconstruction [8, 63, 65, 73] and photorealistic relighting [9, 71], offering a scalable way to train models without the restrictions of real-world data acquisition. Inspired by these previous works, we aim to use synthetic data to improve the model's understanding of the human head and minimize the generalization gap between synthetic data training and real-world inference.

Image or Text to 3D. Generative models have achieved remarkable success in 2D image generation with VAEs [28, 58], GANs [19, 25, 26], and diffusion models [23, 48, 54]. Building on this success, extensive research has extended these models to 3D content generation [18, 41, 43, 66]. Starting with DreamFusion [45], numerous works [36, 46, 51, 57, 64] try to distill NeRF [40, 62, 72] or 3D Gaussians [27] representation from 2D image diffusion using a Score Distillation Sampling (SDS) loss. These methods can produce high-quality results but often encounter challenges such as slow optimization, over-saturated colors, and the Janus problem. To overcome these challenges, recent works [30, 31, 35, 37, 53] generate multi-view images with high consistency, which can be directly used for 3D reconstruction with neural reconstruction methods [27, 40, 62]. However, optimizing NeRF or NeuS remains far from real-time performance. Recent advancements in large 3D reconstruction models (LRMs) [24, 30, 56, 74] offer a pathway to faster 3D reconstruction. Leveraging scalable transformer architectures [15, 59] and large datasets like Objaverse [12, 13], these models effectively capture generalizable 3D priors. Unlike traditional pre-scene optimization methods [27, 40, 62], LRMs employ a feed-forward approach, enabling the prediction of high-quality NeRF, mesh, or 3D Gaussian representations from sparse images in under a second. However, most of these research efforts are applied to general objects, with limited or suboptimal results for 3D head reconstruction.

3. Proposed Method

As shown in Fig. 3, given a single frontal image of a human face y, our goal is to reconstruct a complete 3D head G, represented as 3D Gaussian splats, with detailed facial texture and preserved identity. This requires our system to have prior knowledge on the geometry structure of a human face and the ability to synthesis plausible details which are not visible in the input view. Hence, we train a multi-view diffusion model f_D on synthetic human head data to generate N views $x_0^1, x_0^2, \ldots, x_0^N$ covering 360° of the human head while achieving multi-view consistency and preserving identity. We choose pixel-aligned 3D Gaussians to obtain the final 3D representation. Compared to NeRFs and meshes, 3D Gaussians offer explicit volumetric primitives that better capture subtle facial geometry and fine details. Their semi-transparent kernels naturally model effects like wispy hair and translucency, which are challenging for discrete surfaces or density fields. These generated views $x_0^{1:N}$ from the diffusion model, along with their corresponding Plücker ray coordinates $\mathrm{P}^{(1:N)}$, are input into a transformer-based Gaussian reconstructor f_R to predict a set of 3D Gaussians$G$. Training of the Gaussian reconstructor follows a pre-training process on general objects [12] and a fine-tuning process on synthetic human head data.

3.1. Synthetic Human Head Dataset

We implement a 3D head asset generation pipeline inspired by [65]. Our process begins with a collection of high-quality, artist-created 3D head meshes, which we enhance by incorporating detailed facial components, including eyes, teeth, gums, and both facial and scalp hair. We then augment these base models through rigging for pose variation and blendshape deformation for diverse facial expressions. The final head models are enriched with a set of PBR texture maps, including albedo, normal, roughness, specular, and subsurface scattering maps. At last, we dress the head model with a collection of clothing assets. The entire pipeline is implemented in Blender and the images are rendered with Cycles renderer.

To train our networks, we render images (samples shown in Fig. 4) at 512×512 resolution from 200 unique identities, each with 50 varied appearances, including different hairstyles, skin tones, expressions, clothes, poses, *etc*.

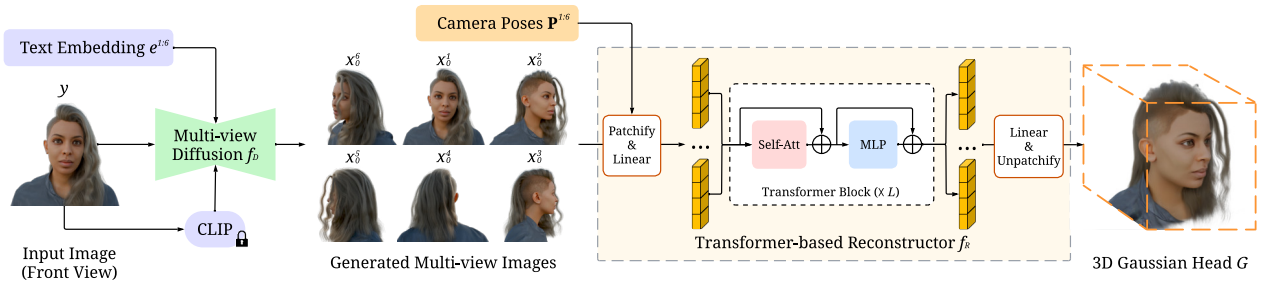

Figure 3. **Overview of *FaceLift*.** Given a single image of a human face y as input, we train an image-conditioned, multi-view diffusion model to generate novel views $x_0^1, \ldots, x_0^N$ covering the entire head. By generating x_0^1 the same as y and leveraging the high-quality synthetic data, our multi-view latent diffusion model can hallucinate unseen views of the human head with high-fidelity and multi-view consistency. We then train a transformer-based reconstructor f_R, which takes multi-view images $x_0^{1:N}$ and their camera poses $\mathbf{P}^{1:N}$ as input and generates 3D Gaussian splats G to represent the human head.

Figure 4. **Synthetic data examples.** Top row: six views for diffusion training. Bottom row: samples of random views for reconstructor training.

We render our training dataset under two types of lighting conditions: (1) ambient light and (2) random HDR environment light. We render six views for each subject to train the multi-view diffusion model. For fine-tuning the transformer-based reconstructor, we render 32 views with random camera poses.

3.2. View Generation

We model the sparse view generation from a single image input as a conditional diffusion process. We use a multi-view diffusion model f_D to generate N views, denoted as $X_0^1, X_0^2, \ldots, X_0^N$, given a single front-facing image y and CLIP text embeddings $e^1, e^2, \ldots, e^N$ corresponding to each generated view. This process is expressed as:

$$\{X_0^1, X_0^2, \ldots, X_0^N\} = f_D(y, \{e^1, e^2, \ldots, e^N\}). \quad (1)$$

We aim to learn the joint distribution of all these views, conditioning on the input image y and text embedding $e^1, e^2, \ldots, e^N$. We denote the joint distribution as:

$$p_{f_D}(x_0^{1:N} \mid y, e^{1:N}) := p_{f_D}(\{x_0^1, \ldots x_0^N\} \mid y, \{e^1, \ldots e^N\}). \quad (2)$$

View Selection. Given a single near frontal view face image with azimuth α, our multi-view diffusion model generates six views with azimuths equal to $\{\alpha, \alpha \pm 45°, \alpha \pm 90°, \alpha + 180°\}$, covering 360 degrees of the human head. Elevation is 0 for all images. We opt for six views as the optimal balance - fewer views compromise detail quality while more views become computationally prohibitive for full head reconstruction. An ablation study comparing four, six, and eight views is presented in Sec. 5.2.

Multi-view Attention. To ensure the consistency of the generated novel views, we use a multi-view attention mechanism to facilitate information propagation and implicitly encode multi-view dependencies. Our attention module encourages multi-view consistency by extending the 2D self-attention mechanism to 3D and enabling interactions across views. Instead of treating each view independently, we apply self-attention across all views simultaneously, allowing information to be shared between them. Specifically, we start with an input tensor of shape $B \times V \times H \times W \times C$, where B is the batch size, V is the number of views, H and W denote the spatial resolution of the intermediate feature maps, and C is the number of feature channels. We reshape this tensor to $B \times VHW \times C$, treating all spatial locations across all views as a unified sequence of tokens for self-attention. This design allows the model to learn multi-view correlations by sharing information across views within the attention layers, enabling it to generate consistent RGB images. We provide an ablation study on the multi-view attention mechanism in the supplementary material.

Input View Reconstruction. During Training, we enforce the first generated view to share the same camera with the input image. In other words, we reconstruct the input view in the view generation process. We find this approach, combined with the multi-view attention mechanism, significantly outperforms the alternative strategy of generating only novel views, which tends to overfit to synthetic training identities and compromises generalization capability as we will show in Sec. 5.1 and the supplementary material.

3.3. Multi-view to 3D Gaussians Reconstruction

Transformer-based Reconstructor. We choose pixel-aligned 3D Gaussians as the final 3D representation. Each Gaussian G_i is defined by position p_i, scale s_i, orientation q_i, opacity α_i, and color features c_i. We use a transformer-based reconstructor f_G to obtain 3D Gaussians from generated multi-view images $x_0^{1:N}$ and their corresponding Plücker ray coordinates [44] $\mathbf{P}^{1:N}$:

$$\{G_i\}_{i=1}^{NHW} = \{p_i, s_i, q_i, \alpha_i, c_i\}_{i=1}^{NHW} = f_G(x_0^{1:N}, \mathbf{P}^{1:N}), \quad (3)$$

Our f_G is a large reconstruction model [24, 74] which follows the implementation of GS-LRM [74]: the N multi-view images are concatenated with their Plücker ray coordinates computed from the camera intrinsic and extrinsic parameters for pose conditioning. Then, the inputs are patchified by dividing the per-view feature map into non-overlapping patches with a patch size of p. Each 2D patch is then flattened into a 1D vector. Finally, a linear layer L is utilized to map the 1D vectors to image patch tokens:

$$\{\mathrm{T}_j^n\}_{j=1,2,\ldots,\frac{HW}{p^2}} = L(\texttt{Patchify}_p(\texttt{Concat}(\mathrm{I}^n, \mathrm{P}^n))). \tag{4}$$

Where $\{\mathrm{T}_j^n\}$ denotes the set of patch tokens for image n, totaling $\frac{HW}{p^2}$ tokens per image. The set of multi-view image tokens $\{\mathrm{T}_j^n\}$ are concatenated and processed through a chain of transformer blocks. Each transformer block is equipped with residual connections [20] and consists of Pre-LayerNorm [2], multi-head Self-Attention [59] and MLP. Later, the output tokens from the transformer are decoded into Gaussian parameters using a single linear layer. Then, the Gaussian parameters are unpatchified into p^2 Gaussians. Finally, we end up with HW Gaussians for each view, where pixel encodes one 3D Gaussian.

Two-stage Training. We find that training the transformer-based reconstructor solely on synthetic human head data leads to inferior texture details when applied to real-world images (see ablation study in Fig. 12). We suspect this limitation arises because the synthetic datasets lack geometric diversity. To address this, we propose a two-stage training approach in which the reconstructor is pre-trained on diverse object data [12] and subsequently fine-tuned using synthetic head data. The pre-training stage enables the reconstructor to learn a broad prior of diverse geometric structures, yielding more detailed and clearer textures in delicate facial regions such as the eyes, nose, and ears. The fine-tuning process then imparts specific knowledge of head structure, producing smoother and more realistic results. During training, we randomly select four input views to reconstruct a total of eight views, four input and four novel views. Following [74], we optimize the model using a combination of MSE and perceptual losses. During inference, the reconstructor processes the six-view outputs from multi-view diffusion model to reconstruct the head.

3.4. Real-world Image Inference

For inference on real-world images, since their intrinsic parameters are unknown, we adopt a camera fov of $50°$, same as during training. To ensure plausible outputs, we first apply an MTCNN face detector to estimate the face's size and center. The image is then resized and cropped/extended to match the average face size and center computed from the training data. We find this alignment compensates for the unknown intrinsic parameters well, ensuring plausible reconstruction results.

4. Experiments

4.1. Experimental Setup

Evaluation Datasets. To quantitatively evaluate *FaceLift*, we establish two benchmarks for single-image to 3D full head reconstruction tasks using publicly available datasets: (1) *Cafca dataset* [8]: We select 40 subjects with 7 to 19 test camera poses each. Since the camera positions are randomly distributed, we manually select the most frontal view as input. Note that this synthetic dataset was independently developed and differs significantly from our training dataset. (2) *Ava-256 dataset* [39]: This studio-captured dataset contains real human subjects. We sample 10 subjects and 10 test camera poses for our evaluation. More details in supplemental. To demonstrate our system's generalization capabilities, we also evaluate on a set of in-the-wild face images for qualitative assessment.

Baselines. We compare our results against three state-of-the-art methods for single-face 3D reconstruction: GGHead [29], PanoHead [1], and Dual Encoder [4]. We perform GAN-inversion to reconstruct 3D head from a given input image using these models. To emphasize the importance of utilizing our synthetic human head data for training, we also compare our method with two methods that focus on general object reconstruction: Era3D [31] and LGM [56]. More comparison results with mesh-based methods are provided in the supplementary material.

We further developed a baseline, *Our MV + LGM*, which leverages the multi-view outputs generated by our diffusion model and employs LGM for reconstruction. This demonstrates that our method can be seamlessly integrated with other reconstruction frameworks to enhance performance on face reconstruction tasks. We tried to fine-tune the LGM reconstructor with our synthetic data, but it provides inferior results with incorrect geometry and artifacts compared with the original weights, which we suspect is due to training data mismatch (see details in the supplementary material).

Evaluation Metrics. We evaluate reconstruction quality using four standard metrics: PSNR, SSIM, LPIPS [75], and DreamSim [16]. To evaluate identity preservation, we perform face verification using ArcFace [14] through the DeepFace [52] implementation.

Implementation Details. Both Cafca [8] and Ava-256 [39] datasets provide multi-view RGB images and corresponding camera poses. However, their camera system differs from the ones utilized in *FaceLift* and baselines. We recalculate the test camera extrinsic in each method's camera system. For a more accurate comparison, we use the Mediapipe facial landmark detection algorithm [38] to identify facial landmarks in both the target images and the rendered outputs, aligning them based on landmark distributions. Details are provided in the supplementary material.

Our system takes approximately 8 seconds to infer a

Method	PSNR ↑	SSIM ↑	LPIPS ↓	DreamSim ↓	ArcFace ↓
GGHead [29]	10.35	0.7406	0.3636	0.3252	0.2681
PanoHead [1]	10.72	0.7594	0.3351	0.2048	0.2183
Dual Encoder [4]	10.78	0.7385	0.3922	0.2785	0.2421
Era3D [31]	13.69	0.7230	0.3662	0.2892	0.2978
LGM [56]	16.52	0.7933	0.3060	0.1552	0.2557
Our MV + LGM [56]	14.13	0.7812	0.2956	0.1282	0.1767
FaceLift	**16.61**	**0.7968**	**0.2694**	**0.1096**	**0.1573**

Table 1. **Quantitative results on Cafca.** *FaceLift* achieves favorable performance on all evaluation metrics.

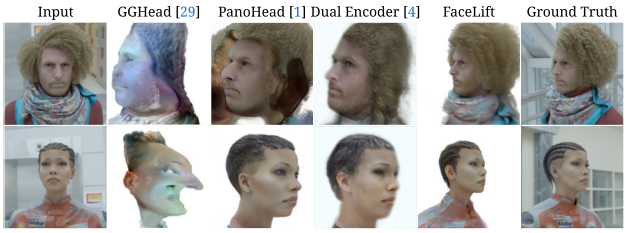

Figure 5. **Visual results on Cafca compared with face reconstruction methods.** *FaceLift* renders novel views that closely match the ground truth, while other baselines often fail to reconstruct the 3D head in correct colors or geometry structures.

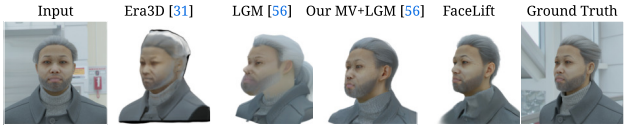

Figure 6. **Visual results on Cafca compared with general objects reconstruction methods.** Comparison with general object reconstruction methods shows the importance of specialized data.

3D Gaussian head from a single image: about 1.5 seconds for preprocessing (background removal, rescaling, etc.), 5.5 seconds for multi-view image generation, and under 1 second for 3D Gaussians reconstruction.

4.2. Experiments on the Cafca Dataset

We report numerical comparison results on Cafca in Tab. 1. *FaceLift* performs favorably against baselines, especially on DreamSim [16] metric, indicating high-quality perceptual similarity. It also achieves better identity preservation performance, as demonstrated by a lower ArcFace [14] distance. We show visual results in Fig. 5 and Fig. 6. *FaceLift* yields rendering results that closely match the ground truth. Compared with other baselines, GGHead [29] does not support full-head rendering, resulting in unrealistic outputs when the view angle significantly deviates from the input. PanoHead [1] struggles with challenging hairstyles, while Dual Encoder [4] produces blurred facial textures. Additionally, Era3D [31] introduces artifacts on the back of the head, and LGM [56] yields inaccurate nose and jaw shapes, underscoring the importance of our synthetic human head data. When integrated with our multi-view diffusion approach, LGM achieves enhanced performance, demonstrating that our method can be seamlessly combined with existing baselines to boost their results.

Method	PSNR ↑	SSIM ↑	LPIPS ↓	DreamSim ↓	ArcFace ↓
Era3D [31]	14.77	0.7963	0.2538	0.2515	0.3721
LGM [56]	14.05	0.8136	0.2476	0.1496	0.3142
Our MV+LGM [56]	15.24	0.8213	0.2292	0.1093	0.2264
FaceLift	**16.52**	**0.8271**	**0.2277**	**0.1065**	**0.1871**

Table 2. **Quantitative results on Ava-256.** *FaceLift* performs favorably than baseline methods in both reconstruction metrics and identity facial identity metric, showing a better generalization ability towards real-captured human images.

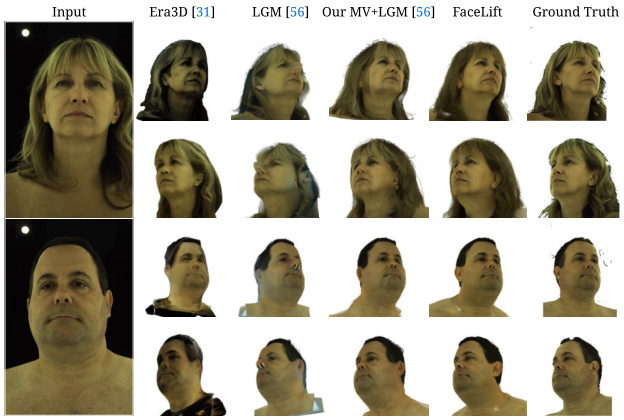

Figure 7. **Visual results on Ava-256.** Compared with baselines, *FaceLift* provides multi-view renderings that are more realistic and similar to ground truth. Era3D fails to deliver delicate facial structures, while LGM generates inaccurate head shapes and colors.

4.3. Experiments on the Ava-256 Dataset

We further evaluate *FaceLift* against other baselines on a studio-captured real human dataset, Ava-256 [50]. GAN-inversion based methods [1, 4, 29] fail to produce reasonable results with the test camera poses in this dataset, so we exclude these baselines. Tab. 2 shows that *FaceLift* outperforms all other baselines across all evaluation metrics, demonstrating superior reconstruction quality and identity preservation. It also highlights *FaceLift*'s strong ability to generalize to real human faces. As shown in Fig. 7, *FaceLift* achieves more realistic head synthesis, while Era3D [31] struggles with accurate skin and hair textures, as well as facial details. LGM [56] produces inaccuracies in the nose shape. When combined with our multi-view diffusion model, LGM yields more accurate geometric structures, yet its texture quality remains inferior to that of *FaceLift*.

4.4. Experiments with In-the-wild Images

We collect in-the-wild human face images and present qualitative results in comparison with other baselines in Fig. 8. Baseline methods often produce undesirable artifacts. For instance, PanoHead [1] frequently fails to render the back of the head and sometimes generates extra eyes at the rear. It also struggles to synthesize hair, shadows, wrinkles, and facial paint accurately, and its outputs lack multi-view consistency (*e.g.*, the girl continues to face the camera in novel view 1 despite a changed pose). Dual Encoder [4] improves

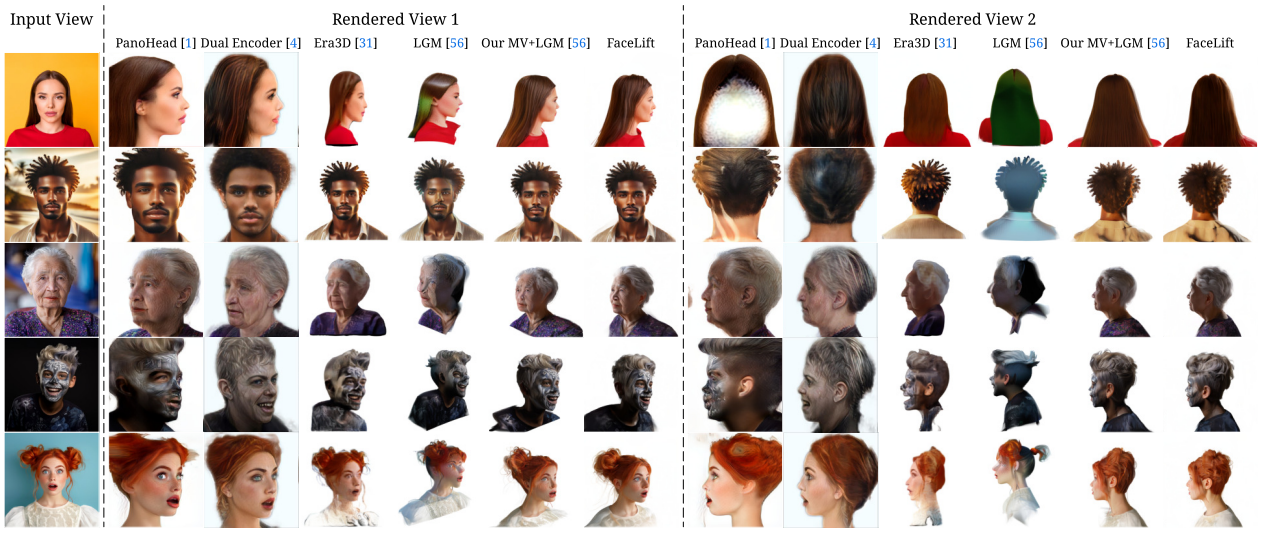

Figure 8. **Visual comparison on in-the-wild data.** *FaceLift* demonstrates great generalization ability and robustness towards in-the-wild images, provides realistic unseen view rendering results. Era3D [31] and LGM [56] generate 3D head representation in inaccurate shape. PanoHead [1] often creates severe artifacts on the back of the head and can not handle challenging hairstyles well. Dual Encoder [4] shows improved performance on reconstructing the back of the head but exhibits more pronounced identity loss.

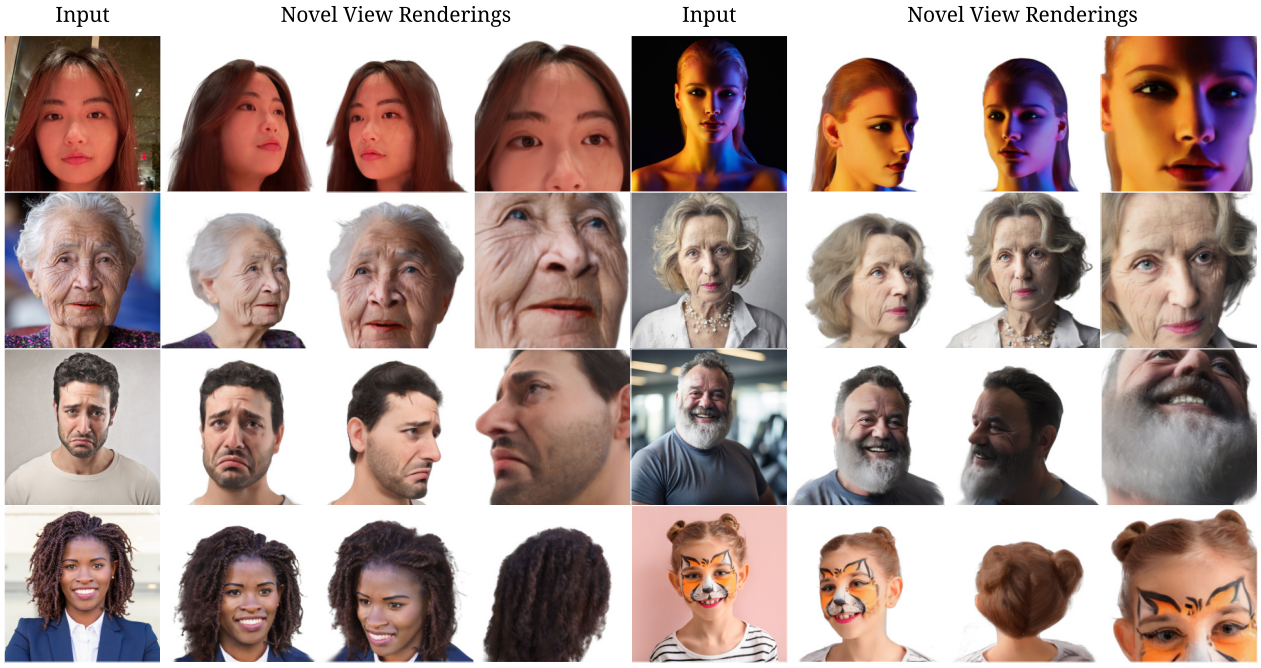

Figure 9. **Results of *FaceLift* on in-the-wild images.** *FaceLift* accurately reconstructs 3D head models under challenging lighting conditions, achieving high fidelity (row 1). It captures fine facial details such as wrinkles (row 2), mustaches (row 3), and individual hairs (row 2 and row 4). Additionally, it remains robust to complex facial expressions (row 3) and various skin tones (row 4). Furthermore, it can realistically reconstruct facial paint (row 4). More results are provided in the supplementary materials.

back-of-head rendering but suffers from severe identity loss (row 2) and fails to accurately reconstruct face paint (row 4). Era3D [31] often produces an inaccurate head shape, particularly from side views, and offers fewer geometric details compared to *FaceLift*. LGM [56] generates Gaussians with inaccurate color and opacity and lacks proper facial geometry, resulting in distorted features. Baseline *Our MV + LGM* shows that our multi-view diffusion model enhances LGM by providing improved facial geometry and texture details. However, the LGM reconstructor still produces Gaussians with inaccurate shapes and opacities, further underscoring the advantages of our transformer-based reconstructor.

We present more *FaceLift*'s novel view rendering results in Fig. 9 to demonstrate *FaceLift*'s ability to produce high-

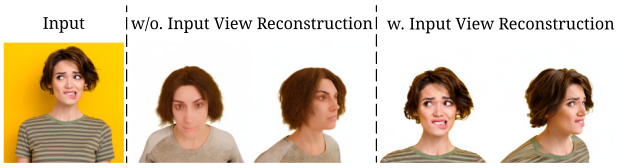

Figure 10. **Importance of input view reconstruction.** The diffusion model that is not trained to perform the input view reconstruction, *i.e.*, *w/o. Input View Reconstruction*, overfits to synthetic training distribution, suffers from severe identity loss during inference. Trained with input view reconstruction, our method preserves the input identity and expression faithfully.

fidelity, realistic 3D head reconstructions with intricate details across a variety of challenging scenarios. *FaceLift* effectively handles faces under various lighting conditions. It can especially render realistic novel view images given a photo captured with an iPhone under dark lighting conditions (row 1 column 1), emphasizing its robustness and potential for real-world application. It reconstructs facial details with high fidelity, especially the wrinkles and folds on the face caused by extreme expression. *FaceLift* also excels at reconstructing challenging textures, such as mustaches and hair. Furthermore, it faithfully reconstructs facial paint, despite such data not being included in our synthetic face dataset, showcasing its strong generalization ability.

5. Ablation Study

5.1. Input View Reconstruction

We conduct an ablation study to demonstrate the importance of reconstructing the input view during training. For comparison, we train a multi-view diffusion model that generates six novel views. In this baseline, the first generated view's elevation is adjusted from $0°$ to $20°$, while the remaining views adopt the same camera poses as in our default setting. We refer to this variant as *w/o. Input View Reconstruction*. Fig. 10 presents the view generation results of the two diffusion models when applied to real-world images. Without the input view reconstruction task, the model trained on the synthetic dataset generates views within a limited distribution, leading to noticeable identity loss. Moreover, it loses its ability to preserve facial expressions and face paint. In contrast, incorporating the input view reconstruction task during training enables our diffusion model to faithfully regenerate the input view, significantly improving its generalization ability. Quantitative comparison is provided in the supplementary material.

5.2. Number of Views

We evaluate three configurations: four views (front, left, right, and back), six views (adding front-left and front-right), and eight views (further including front-top and front-bottom). Fig. 11 compares the baselines using different numbers of input views. With only four views, the reconstructor fails to capture a complete forehead; however,

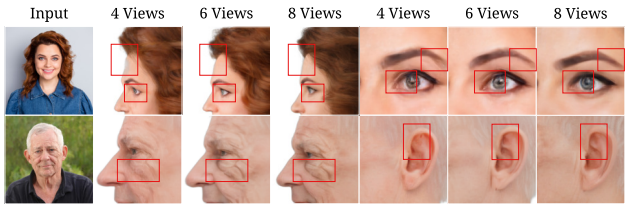

Figure 11. **Number of input views of Gaussian reconstructor.** Using six views strike a good balance between reconstruction quality and computational efficiency.

Figure 12. **Two-stage training of reconstructor.** Without pre-training on general objects, the reconstructor fails to produce clear textures in the reconstruction results. Meanwhile, without fine-tuning on synthetic human head data, the model lacks a refined understanding of facial structures, including the eyes and nose.

with six views, it reconstructs the eyes and eyebrows more smoothly and renders challenging textures—such as facial wrinkles and ear folds—more realistically. Eight views do not offer significant visual improvements, and incur a higher computational cost in both stages. Thus, we find that six views achieve a good balance between reconstruction quality and computational efficiency.

5.3. Two-stage Reconstructor Training

As illustrated in Sec. 3.3, our Gaussian reconstructor follows a two-stage training pipeline. Fig. 12 shows that pre-training on general objects helps the model learn a diverse prior of geometric structures, resulting in clearer textures on delicate facial regions. Meanwhile, fine-tuning on synthetic human head data enables the reconstructor to gain a more refined understanding of facial structure, thereby enhancing the accuracy of features such as the eyes, nose, and hair.

6. Conclusions

We propose *FaceLift*, a feed-forward approach that lifts a single facial image to a detailed 3D reconstruction with preserved identity features. Our method uses multi-view diffusion to generate unobservable views and employs a transformer-based reconstructor to reconstruct 3D Gaussian splats, enabling high-quality novel view synthesis. To overcome the difficulty of capturing real-world multi-view human head images, we render high-quality synthetic data for training and show that, despite being trained solely on synthetic data, *FaceLift* can reconstruct 3D heads from real-world captured images with high fidelity. Compared with baselines [1, 4, 29, 31, 56], *FaceLift* generates 3D head representation with finer geometry and texture details and exhibits better identity preservation ability.

7. Acknowledgement

This was supported in part by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korean Government (MSIT) (No. RS-2024-00457882, National AI Research Lab Project).

We appreciate the insightful discussions with Kai Zhang, Hao Tan, Zexiang Xu, Sai Bi, Sumit Chaturvedi, Hanwen Jiang, Yu-Ju Tsai, Kuan-Chih Huang, Chengxu Liu, and Dingyi Dai. We thank Nathan Carr and Kalyan Sunkavalli for their support.

References

[1] Sizhe An, Hongyi Xu, Yichun Shi, Guoxian Song, Umit Y Ogras, and Linjie Luo. PanoHead: Geometry-aware 3D full-head synthesis in 360°. In *CVPR*, 2023. 1, 3, 5, 6, 7, 8

[2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 5

[3] Ananta R. Bhattarai, Matthias Nießner, and Artem Sevastopolsky. TriPlaneNet: An encoder for EG3D inversion. In *WACV*, 2024. 2

[4] Bahri Batuhan Bilecen, Ahmet Berke Gokmen, and Aysegul Dundar. Dual encoder GAN inversion for high-fidelity 3D head reconstruction from single images. In *NeurIPS*, 2024. 5, 6, 7, 8

[5] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3D faces. In *SIGGRAPH*, 1999. 2

[6] James Booth, Anastasios Roussos, Stefanos Zafeiriou, Allan Ponniah, and David Dunaway. A 3D morphable model learnt from 10,000 faces. In *CVPR*, 2016. 2

[7] Marcel C Bühler, Kripasindhu Sarkar, Tanmay Shah, Gengyan Li, Daoye Wang, Leonhard Helminger, Sergio Orts-Escolano, Dmitry Lagun, Otmar Hilliges, Thabo Beeler, et al. Preface: A data-driven volumetric prior for few-shot ultra high-resolution face synthesis. In *ICCV*, 2023. 3

[8] Marcel C Bühler, Gengyan Li, Erroll Wood, Leonhard Helminger, Xu Chen, Tanmay Shah, Daoye Wang, Stephan Garbin, Sergio Orts-Escolano, Otmar Hilliges, et al. Cafca: High-quality novel view synthesis of expressive faces from casual few-shot captures. In *SIGGRAPH Asia*, 2024. 2, 3, 5, 12, 15, 17

[9] Sumit Chaturvedi, Mengwei Ren, Yannick Hold-Geoffroy, Jingyuan Liu, Julie Dorsey, and Zhixin Shu. SynthLight: Portrait relighting with diffusion model by learning to re-render synthetic faces. In *CVPR*, 2025. 3

[10] Xiyi Chen, Marko Mihajlovic, Shaofei Wang, Sergey Prokudin, and Siyu Tang. Morphable diffusion: 3d-consistent diffusion for single-image avatar creation. In *CVPR*, 2024. 3

[11] Antonia Creswell and Anil Anthony Bharath. Inverting the generator of a generative adversarial network. In *TNNLS*, 2018. 3

[12] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3D objects. In *CVPR*, 2023. 2, 3, 5

[13] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, et al. Objaverse-XL: A universe of 10M+ 3D objects. In *NeurIPS*, 2024. 3

[14] Jiankang Deng, Jia Guo, Jing Yang, Niannan Xue, Irene Kotsia, and Stefanos Zafeiriou. ArcFace: Additive angular margin loss for deep face recognition. In *IEEE TPAMI*, 2022. 5, 6

[15] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2020. 3

[16] Stephanie Fu, Netanel Tamir, Shobhita Sundaram, Lucy Chai, Richard Zhang, Tali Dekel, and Phillip Isola. DreamSim: Learning new dimensions of human visual similarity using synthetic data. In *NeurIPS*, 2023. 5, 6

[17] Chen Gao, Yichang Shih, Wei-Sheng Lai, Chia-Kai Liang, and Jia-Bin Huang. Portrait neural radiance fields from a single image. *arXiv preprint arXiv:2012.05903*, 2020. 3

[18] Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. GET3D: A generative model of high quality 3D textured shapes learned from images. In *NeurIPS*, 2022. 3

[19] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014. 1, 2, 3

[20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5

[21] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. 16

[22] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 15

[23] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 1, 3

[24] Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. LRM: Large reconstruction model for single image to 3D. In *ICLR*, 2024. 3, 5

[25] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 1, 2, 3

[26] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, 2020. 2, 3

[27] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D gaussian splatting for real-time radiance field rendering. In *ACM TOG*, 2023. 1, 3

[28] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014. 3

[29] Tobias Kirschstein, Simon Giebenhain, Jiapeng Tang, Markos Georgopoulos, and Matthias Nießner. GGHead: Fast and generalizable 3D gaussian heads. In *SIGGRAPH Asia*, 2024. 5, 6, 8

[30] Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan Sunkavalli, Greg Shakhnarovich, and Sai Bi. Instant3D: Fast text-to-3D with sparse-view generation and large reconstruction model. In *ICLR*, 2024. 3

[31] Peng Li, Yuan Liu, Xiaoxiao Long, Feihu Zhang, Cheng Lin, Mengfei Li, Xingqun Qi, Shanghang Zhang, Wenhan Luo, Ping Tan, et al. Era3D: High-resolution multiview diffusion using efficient row-wise attention. In *NeurIPS*, 2024. 3, 5, 6, 7, 8

[32] Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. In *ACM TOG*, 2017. 1, 2

[33] Xueting Li, Shalini De Mello, Sifei Liu, Koki Nagano, Umar Iqbal, and Jan Kautz. Generalizable one-shot 3D neural head avatar. In *NeurIPS*, 2024. 2

[34] Jiangke Lin, Yi Yuan, Tianjia Shao, and Kun Zhou. Towards high-fidelity 3D face reconstruction from in-the-wild images using graph convolutional networks. In *CVPR*, 2020. 2

[35] Minghua Liu, Chao Xu, Haian 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. In *NeurIPS*, 2023. 3

[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. In *ICCV*, 2023. 3

[37] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3D: Single image to 3D using cross-domain diffusion. In *CVPR*, 2024. 3

[38] Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Guang Yong, Juhyun Lee, et al. Mediapipe: A framework for building perception pipelines. *arXiv preprint arXiv:1906.08172*, 2019. 5

[39] Julieta Martinez, Emily Kim, Javier Romero, Timur Bagautdinov, Shunsuke Saito, Shoou-I Yu, Stuart Anderson, Michael Zollhöfer, Te-Li Wang, Shaojie Bai, Chenghui Li, Shih-En Wei, Rohan Joshi, Wyatt Borsos, Tomas Simon, Jason Saragih, Paul Theodosis, Alexander Greene, Anjani Josyula, Silvio Mano Maeta, Andrew I. Jewett, Simon Venshtain, Christopher Heilman, Yueh-Tung Chen, Sidi Fu, Mohamed Ezzeldin A. Elshaer, Tingfang Du, Longhua Wu, Shen-Chi Chen, Kai Kang, Michael Wu, Youssef Emad, Steven Longay, Ashley Brewer, Hitesh Shah, James Booth, Taylor Koska, Kayla Haidle, Matt Andromalos, Joanna Hsu, Thomas Dauer, Peter Selednik, Tim Godisart, Scott Ardisson, Matthew Cipperly, Ben Humberston, Lon Farr, Bob Hansen, Peihong Guo, Dave Braun, Steven Krenn, He Wen, Lucas Evans, Natalia Fadeeva, Matthew Stewart, Gabriel Schwartz, Divam Gupta, Gyeongsik Moon, Kaiwen Guo, Yuan Dong, Yichen Xu, Takaaki Shiratori, Fabian Prada, Bernardo R. Pires, Bo Peng, Julia Buffalini, Autumn Trimble, Kevyn McPhail, Melissa Schoeller, and Yaser Sheikh. Codec Avatar Studio: Paired human captures for complete, driveable, and generalizable avatars. In *NeurIPS*, 2024. 2, 3, 5, 15, 17

[40] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 1, 3

[41] Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Joon Park, and Ira Kemelmacher-Shlizerman. StyleSDF: High-resolution 3D-consistent image and geometry generation. In *CVPR*, 2022. 3

[42] Foivos Paraperas Papantoniou, Alexandros Lattas, Stylianos Moschoglou, and Stefanos Zafeiriou. Relightify: Relightable 3D faces from a single image via diffusion models. In *ICCV*, 2023. 3

[43] Dario Pavllo, Graham Spinks, Thomas Hofmann, Marie-Francine Moens, and Aurélien Lucchi. Convolutional generation of textured 3D meshes. In *NeurIPS*, 2020. 3

[44] Julius Plucker. Xvii. on a new geometry of space. *Philosophical Transactions of the Royal Society of London*, 1865. 4

[45] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. DreamFusion: Text-to-3D using 2D diffusion. In *ICLR*, 2023. 3

[46] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skorokhodov, Peter Wonka, Sergey Tulyakov, et al. Magic123: One image to high-quality 3D object generation using both 2D and 3D diffusion priors. In *ICLR*, 2024. 3

[47] Elad Richardson, Matan Sela, and Ron Kimmel. 3D face reconstruction by learning from synthetic data. In *3DV*, 2016. 2

[48] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2, 3

[49] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 16

[50] Shunsuke Saito, Gabriel Schwartz, Tomas Simon, Junxuan Li, and Giljoo Nam. Relightable gaussian codec avatars. In *CVPR*, 2024. 6

[51] Kyle Sargent, Zizhang Li, Tanmay Shah, Charles Herrmann, Hong-Xing Yu, Yunzhi Zhang, Eric Ryan Chan, Dmitry Lagun, Li Fei-Fei, Deqing Sun, et al. ZeroNVS: Zero-shot 360-degree view synthesis from a single real image. In *CVPR*, 2024. 3

[52] Sefik Ilkin Serengil and Alper Ozpinar. HyperExtended LightFace: A facial attribute analysis framework. In *ICEET*, 2021. 5

[53] Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen, Chong Zeng, and Hao Su. Zero123++: A single image to consistent multi-view diffusion base model. *arXiv preprint arXiv:2310.15110*, 2023. 3

[54] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021. 3, 16

[55] Wenqiang Sun, Shuo Chen, Fangfu Liu, Zilong Chen, Yueqi Duan, Jun Zhang, and Yikai Wang. DimensionX: Create any 3D and 4D scenes from a single image with controllable video diffusion. *arXiv preprint arXiv:2411.04928*, 2024. 12

[56] Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. LGM: Large multi-view gaussian model for high-resolution 3D content creation. In *ECCV*, 2024. 3, 5, 6, 7, 8, 17, 18

[57] Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. DreamGaussian: Generative gaussian splatting for efficient 3D content creation. In *ICLR*, 2024. 3

[58] Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In *NeurIPS*, 2017. 3

[59] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 3, 5

[60] Thomas Vetter and Volker Blanz. Estimating coloured 3D face models from single images: An example based approach. In *ECCV*, 1998. 1, 2

[61] Vishal Vinod, Tanmay Shah, and Dmitry Lagun. TEGLO: High fidelity canonical texture mapping from single-view images. In *WACV*, 2024. 3

[62] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. NeuS: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In *NeurIPS*, 2021. 3

[63] Tengfei Wang, Bo Zhang, Ting Zhang, Shuyang Gu, Jianmin Bao, Tadas Baltrusaitis, Jingjing Shen, Dong Chen, Fang Wen, Qifeng Chen, et al. Rodin: A generative model for sculpting 3D digital avatars using diffusion. In *CVPR*, 2023. 3

[64] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. ProlificDreamer: High-fidelity and diverse text-to-3D generation with variational score distillation. In *NeurIPS*, 2024. 3

[65] Erroll Wood, Tadas Baltrušaitis, Charlie Hewitt, Sebastian Dziadzio, Thomas J Cashman, and Jamie Shotton. Fake it till you make it: Face analysis in the wild using synthetic data alone. In *ICCV*, 2021. 3

[66] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling. In *NeurIPS*, 2016. 3

[67] Kailu Wu, Fangfu Liu, Zhihan Cai, Runjie Yan, Hanyang Wang, Yating Hu, Yueqi Duan, and Kaisheng Ma. Unique3d: High-quality and efficient 3d mesh generation from a single image. *arXiv preprint arXiv:2405.20343*, 2024. 12, 13

[68] Yiqian Wu, Jing Zhang, Hongbo Fu, and Xiaogang Jin. LPFF: A portrait dataset for face generators across large poses. In *ICCV*, 2023. 1

[69] Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. Structured 3d latents for scalable and versatile 3d generation. In *CVPR*, 2024. 12, 13

[70] Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. *arXiv preprint arXiv:2404.07191*, 2024. 12, 13

[71] Yu-Ying Yeh, Koki Nagano, Sameh Khamis, Jan Kautz, Ming-Yu Liu, and Ting-Chun Wang. Learning to relight portrait images via a virtual light stage and synthetic-to-real adaptation. In *ACM TOG*, 2022. 3

[72] Ziyang Yuan, Yiming Zhu, Yu Li, Hongyu Liu, and Chun Yuan. Make encoder great again in 3D GAN inversion through geometry and occlusion-aware encoding. In *ICCV*, 2023. 1, 2, 3

[73] Bowen Zhang, Yiji Cheng, Chunyu Wang, Ting Zhang, Jiaolong Yang, Yansong Tang, Feng Zhao, Dong Chen, and Baining Guo. RodinHD: High-fidelity 3D avatar generation with diffusion models. In *ECCV*, 2024. 2, 3

[74] Kai Zhang, Sai Bi, Hao Tan, Yuanbo Xiangli, Nanxuan Zhao, Kalyan Sunkavalli, and Zexiang Xu. GS-LRM: Large reconstruction model for 3D gaussian splatting. In *ECCV*, 2024. 3, 5, 17

[75] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 5

[76] Xiaozheng Zheng, Chao Wen, Zhaohu Li, Weiyi Zhang, Zhuo Su, Xu Chang, Yang Zhao, Zheng Lv, Xiaoyuan Zhang, Yongjie Zhang, Guidong Wang, and Xu Lan. Headgap: Few-shot 3d head avatar via generalizable gaussian priors. In *3DV*, 2025. 3