

Large-Scale High-Quality 3D Gaussian Head Reconstruction from Multi-View Captures

Evangelos Ntavelis, Sean Wu[†], Mohamad Shahbazi, Fabio Maninchedda, Dmitry Kostiaev, Artem Sevastopolsky, Vittorio Megaro, Trevor Phillips, Alejandro Blumentals, Shridhar Ravikumar, Mehak Gupta, Reinhard Knothe, Jeronimo Bayer, Matthias Vestner, Simon Schaefer, Thomas Etterlin, Christian Zimmermann, Alexey Artemov, Mathias Deschler, Peter Kaufmann, Stefan Brugger, Sebastian Martin, Brian Amberg, and Tom Runia

Apple

<https://apple.github.io/ml-headsup/>

Abstract. We propose *HeadsUp*, a scalable feed-forward method for reconstructing high-quality 3D Gaussian heads from large-scale multi-camera setups. Our method employs an efficient encoder-decoder architecture that compresses input views into a compact latent representation. This latent representation is then decoded into a set of UV-parameterized 3D Gaussians anchored to a neutral head template. This UV representation decouples the number of 3D Gaussians from the number and resolution of input images, enabling training with many high-resolution input views. We train and evaluate our model on an internal dataset with more than 10 000 subjects, which is an order of magnitude larger than existing multi-view human head datasets. *HeadsUp* achieves state-of-the-art reconstruction quality and generalizes to novel identities without test-time optimization. We extensively analyze the scaling behavior of our model across identities, views, and model capacity, revealing practical insights for quality-compute trade-offs. Finally, we highlight the strength of our latent space by showcasing two downstream applications: generating novel 3D identities and animating the 3D heads with expression blendshapes.

1 Introduction

High-fidelity 3D head assets are foundational to photorealistic digital humans, enabling convincing and authentic social co-presence in immersive digital environments. Such assets are particularly valuable for view-consistent, close-up rendering in applications such as telepresence, actor digitization, and content creation [36]. To achieve this level of quality, multi-camera capture systems have become a standard approach for producing dense calibrated images of human heads [11]. However, it remains a challenge to reliably and efficiently transform these detailed captures into compact 3D reconstructions at scale [34, 76].

[†]Work done during internship at Apple.



Fig. 1. We introduce *HeadsUp*, a novel feed-forward approach leveraging 3D Gaussians to predict high-quality avatars. By scaling to thousands of subjects and diverse expressions, our method achieves exceptional rendering quality on completely held-out subjects. Notice the accurate, high-resolution recovery of intricate fine details, such as eyelashes, complex earrings, teeth and tongue. The figure displays renders from novel subjects unseen during training.

Existing solutions for converting multi-view head captures into renderable assets span a spectrum that reflects a recurring tension between reconstruction fidelity and throughput. At one end, instance-specific optimization methods, such as per-subject fitting using Neural Radiance Fields [4, 33, 48] or 3D Gaussian Splatting (3DGS) [27, 58], serve as strong baselines for high-quality view-consistent rendering. However, the computational cost of per-identity optimization makes their large-scale deployment challenging. At the other end, recent feed-forward reconstruction methods [6, 8, 34, 85] amortize computation across datasets and enable fast inference, but their compute and memory costs typically scale with the number and resolution of input views. This hinders the full utilization of dense, high-resolution capture setups. Orthogonally, another line of work targets animatable avatars by predicting geometry and appearance in a canonical, model-conditioned space [9, 28, 76]. While this supports temporal coherence and control, it often sacrifices the representation capacity needed for high-quality rendering. Together, these trends suggest a clear need for methods that can fully leverage multi-camera rigs (i.e., many views, high resolution, many identities) while producing compact head assets suitable for photorealistic rendering. To address this need, we prioritize high-fidelity reconstruction quality from high-resolution multi-view captures spanning thousands of identities and millions of frames.

We propose *HeadsUp*, a scalable feed-forward approach for reconstructing high-fidelity 3D Gaussian head assets from large-scale multi-camera studio captures. Specifically, our method transforms input multi-view images to the UV parameterization of a set of 3D Gaussians attached to a neutral head template. In contrast to pixel-aligned prediction [23, 34], the UV-space formulation decouples the number of output Gaussians from the number and resolution of input images. This allows our method to scale gracefully with the number of high-resolution input views, which is crucial for resolving fine details such as hair strands and jewelry. Our feed-forward model is based on an encoder-decoder architecture: a cross-attention transformer efficiently encodes the input images into a compact 2D

latent representation, which is then decoded into the corresponding Gaussian UV maps. To better capture the high-frequency details, we introduce a background model to improve fine boundary structures such as hair, and a high-resolution finetuning stage that leverages multi-scale and region-specific losses to enhance overall render quality and the fidelity of core facial features like the eyes and mouth.

To demonstrate the reconstruction quality and scalability of our method, we train and evaluate our model on a large-scale internal dataset that is an order of magnitude larger than existing multi-view face datasets, comprising over 10 000 subjects with diverse appearances and facial expressions. For comparison with prior work, we also report results on the publicly available Ava-256 dataset [47]. *HeadsUp* achieves state-of-the-art reconstruction quality on novel identities and expressions (examples in Fig. 1). We further characterize the scaling behavior of our model with respect to the number of training subjects and input views, as well as the model’s capacity. Finally, we highlight the strength of our learned latent space by showcasing two downstream applications: generating novel 3D identities with latent diffusion models and animating the 3D heads from expression blendshapes.

In summary, our main contributions are as follows:

- We address the challenge of large-scale, high-fidelity 3D head reconstruction from multi-camera studio captures spanning thousands of identities and millions of frames.
- We introduce *HeadsUp*, a feed-forward model that predicts UV maps that encode the parameters of 3D Gaussians anchored to a neutral head template. Our model decouples the output representation from the input resolution and view count, enabling efficient scaling on dense, high-resolution data.
- We introduce explicit background modeling and a tailored training strategy to improve reconstruction in high-frequency regions (hair, eyes, and mouth).
- We achieve state-of-the-art reconstruction quality on novel identities and expressions, and characterize the scaling behavior of our model across identities, views, and model capacity.

2 Related Work

Novel View Synthesis of Human Heads. Various 3D representations have been explored for head reconstruction and novel view synthesis. Earlier attempts based on Neural Radiance Fields [2, 16, 33, 51, 89] have succeeded at photorealistic rendering but required lengthy per-scene optimization. Implicit surface methods [87] and point-based approaches [84, 88] offered alternatives with different trade-offs between quality and efficiency. Hybrid methods like MonoNPHM [17] combine neural parametric models with implicit representations for monocular reconstruction. Works on Codec Avatars [38, 43, 46, 58] and others [66] achieve high-fidelity head avatars through learning on multiple dense multi-view captures. Often, these methods involve per-subject optimization or fine-tuning at test

time. 3DGS optimization-based approaches have quickly become popular due to their efficiency and more explicit representation, albeit limited in scalability because of the large number of Gaussians required. GaussianAvatars [54] ties 3D Gaussians to a parametric model [41] for fully controllable heads but requires minutes of optimization per identity, excluding multi-view tracking time. SplatttingAvatar [59] embeds Gaussians within triangle meshes for hybrid mesh-Gaussian avatars. GAF [64] distills multi-view diffusion priors using pseudo-ground truths to enhance monocular avatar reconstruction. These works, along with [65, 71] and others, demonstrate that 3DGS is an excellent representation for detailed head avatars. However, the reliance on slow optimization makes scaling to thousands of subjects with high-resolution studio captures challenging.

Feed-forward 3D Head Reconstruction. Following the advances in Large Reconstruction Models (LRMs) [6, 8, 21, 63, 69, 72, 81], recent works develop feed-forward 3DGS head models to bypass per-subject optimization.

Multi-view feed-forward methods address the well-constrained setting of reconstruction from multiple input images. GPAvatar [10] learns efficient Gaussian projections from multi-view inputs by employing deep feature extractors. HeadGAP [86] learns generalizable Gaussian priors from multi-view data and personalizes to new identities from a few images, though it still requires a per-subject adaptation step. Avat3r [34] regresses animatable 3D head avatars from as few as four images using DUST3R position maps [70] and Sapiens features [29]. While template-free, its pixel-aligned Gaussians scale linearly with input views, causing memory constraints and potential temporal inconsistencies. The reliance on DUST3R and Sapiens makes the pipeline time- and memory-constrained in a practical setting. Similarly, FastGHA [23] achieves few-shot real-time animation via pixel-aligned features, yet faces comparable scaling limitations. The concurrent work MATCH [53] efficiently predicts Gaussian textures within the fixed UV layout of a template mesh, whereas our method aggregates multi-view features into a compact latent representation prior to UV decoding. Finally, Pippo [26] achieves high quality results by employing a diffusion transformer trained on a large-scale dataset. While practical for single images, inference scales to minutes for multi-view captures, and strict 3D view consistency is not guaranteed. Parallel works on full-body reconstruction explore similar concepts, such as efficiently reconstructing animatable humans from pose-free images [55, 56], regressing Gaussians in UV space [35], or training DiT-based generators [75]. Despite achieving impressive full-body results, these methods remain bottlenecked by the limited facial resolution of available datasets.

Single-view feed-forward methods tackle the ill-posed problem of reconstructing human heads from a single input image. GAGAvatar [9] is one of the first generalizable triplane-based Gaussian head models. Portrait4D [13, 14] learns novel view synthesis and driving from synthetic 2D data and pseudo multi-view supervision, generated from a triplane-based generator. However, triplane representations often struggle with extreme novel viewpoints and are sensitive to image cropping. LAM [20] and PanoLAM [39] perform one-shot head synthesis based on UV-aligned Gaussians, similar to our method but with a focus on rigging.

FastAvatar [42] predicts residual Gaussians from a canonical template in <10 ms, specifically targeting the fast rigging setting. PercHead [50] uses perceptual supervision from DINOv2 [49] and SAM [31] for disentangled geometry and appearance control. Our method partially borrows inspiration from some of these methods and uses a UV-aligned Gaussian map that naturally supports an arbitrary number of views, including a single view, while remaining fast and supporting a wide range of novel viewpoints. Rather than optimizing for real-time rigging, in this work we prioritize maximum reconstruction quality. Like Pippo [26], several recent methods rely on diffusion priors to reconstruct human heads [7, 40, 44, 45, 73, 74, 78–80], but this reliance comes at computational cost and compromises rendering fidelity.

3D GAN inversion represents a parallel track for single-image 3D head synthesis. PanoHead [1] and SphereHead [37] extend EG3D [5] to 360° head synthesis with back-of-head supervision. Encoder-based latent inversion methods like TriPlaneNet [3] and GOAE [77], alongside optimization-based PTI [57], facilitate reconstruction from real images, while DiffPortrait3D [19] and VOODOO3D [67] focus on one-shot reenactment. InvertAvatar [83] adapts this inversion strategy to multiple images for incremental improvement. However, triplanes can fundamentally limit the level of achievable resolution and expressivity. Furthermore, triplane-based generators rely on external 2D head pose estimation that can introduce compounding errors; these methods also usually impose strict cropping requirements on the input image and camera assumptions [5, 61].

Summary. Ultimately, our work distinguishes itself from relevant prior work by combining several key properties: (1) a fully feed-forward architecture that avoids lengthy optimization to a novel subject and decouples memory from the input view count, enabling massive scalability; (2) universal support for monocular, sparse, or dense multi-view inputs; (3) independence from strict parametric face models, allowing for natural reconstruction of accessories and diverse expressions; and (4) state-of-the-art photorealism and subject recognizability in novel views.

3 Method

Here we introduce our method *HedsUp*. Fig. 2 shows the overall architecture of our model. Given N calibrated time-synchronized images, our model jointly predicts 3D Gaussians for the foreground and background. Explicit background modeling bypasses the need for foreground matting, which often fails to accurately segment high-frequency boundary details like hair or jewelry.

Our model is based on UV-parameterized 3D Gaussians anchored to a template mesh. Our feed-forward network consists of two main components: a *multi-view encoder* based on a transformer architecture to extract latents from the input images, and a *3D Gaussian decoder* to map the latents to the 3D Gaussians’ UV parameters. To ensure high-fidelity rendering of core facial features without prohibitive computational costs, we employ a two-stage training strategy where a high-resolution finetuning stage leverages region-specific and multi-scale supervision. Below, we detail our model architecture and training methodology.




Fig. 2. Overview of *HeadsUp*. Our method reconstructs high-fidelity 3D Gaussian heads from multi-view images. Given a set of input views, our model utilizes a transformer-based encoder and a 3D Gaussian decoder to predict UV-parameterized 3D Gaussians for both the foreground and background. The model is trained end-to-end using a combination of photometric and perceptual supervision.

3.1 UV-Parameterized 3D Gaussian Splatting

Our model predicts a multi-channel 2D UV map, representing a fixed number of Gaussians anchored to a template mesh. The channels of the UV map correspond to the Gaussian attributes: position $\mu \in \mathbb{R}^3$, scale $s \in \mathbb{R}_+^3$, quaternion rotation $\mathbf{q} \in \mathbb{R}^4$, opacity $\alpha \in [0, 1]$, and spherical harmonic color coefficients $\{\mathbf{c}^{(\ell, m)}\}$ up to degree $L = 1$. The foreground Gaussians are anchored to a neutral head template shared among all identities in a canonical coordinate system. The head canonical coordinate system is defined with its origin at the mid-pupil point of the template and its orientation aligned to the Frankfurt plane [22]. The background Gaussians are anchored to a sphere template fitted to the capture rig and are transformed into the canonical coordinate system for joint rendering.

The UV formulation in our method offers several advantages, including:

- (1) *Shared geometric prior.* The canonical mesh topology provides consistent spatial structure across subjects and expressions, allowing the network to focus on appearance and local variations rather than coarse global head geometry.
- (2) *Efficient multi-view aggregation.* The mesh-anchored UV parameterization decouples the number of output Gaussians from the number and resolution of input images, allowing for efficient aggregation of information from dense and high-resolution captures.
- (3) *Robustness to tracking errors.* As the 3D Gaussians are anchored to a fixed neutral template, our representation only requires rigid head pose tracking and does not rely on fragile and error-prone facial expression tracking for inference.

3.2 Network Architecture

Here we describe our network architecture, composed of a *multi-view encoder* and a *Gaussian UV decoder*.

Multi-View Encoder. Given N calibrated input views $\{\mathbf{I}_i\}_{i=1}^N$, $\mathbf{I}_i \in \mathbb{R}^{3 \times H \times W}$, and their camera extrinsics $\{\mathbf{E}_i\}_{i=1}^N$, $\mathbf{E}_i \in \mathbb{R}^{4 \times 4}$, along with intrinsics $\{\mathbf{K}_i\}_{i=1}^N$, $\mathbf{K}_i \in \mathbb{R}^{3 \times 3}$, we first patchify the images and convert each into patch embeddings $\mathbf{e}_p^i \in \mathbb{R}^{d \times h \times w}$. To explicitly encode the camera geometry into the features, we concatenate per-patch Plücker embeddings [60] with the corresponding patch embeddings along the channel dimension:

$$\mathbf{f}_p^i = \text{Concat}(\mathbf{e}_p^i, \mathcal{P}(\mathbf{K}_i, \mathbf{E}_i)) \in \mathbb{R}^{(d+6) \times h \times w} \quad (1)$$

where $\mathcal{P}(\cdot)$ computes the 6D Plücker coordinates. Two parallel convolutional encoders process these features to disentangle each input view into foreground and background features, $\mathbf{F}_i^{\text{fg}} \in \mathbb{R}^{c \times h_f \times w_f}$ and $\mathbf{F}_i^{\text{bg}} \in \mathbb{R}^{c' \times h_f \times w_f}$, respectively.

We employ a transformer architecture [15, 68] to map these unstructured multi-view foreground features to a low-resolution 2D latent representation of the target 3D head. The transformer converts a 2D grid of learnable tokens $\mathbf{Q} \in \mathbb{R}^{d_z \times h_z \times w_z}$ to the 2D latent $\mathbf{Z} \in \mathbb{R}^{d_z \times h_z \times w_z}$ by aggregating information from multi-view features through cross-attention layers:

$$\mathbf{Z} = \text{CrossAttnTransformer}(\mathbf{Q}, \mathbf{F}_{\text{fg}}). \quad (2)$$

When discussing our latent representation throughout the paper, we refer to the foreground latent $\mathbf{Z}$.

To model the background, we use a shallow convolutional network with global average pooling to map the stacked multi-view background features $\{\mathbf{F}_i^{\text{bg}}\}_{i=1}^N$ to the background latent $\mathbf{z}_{\text{bg}} \in \mathbb{R}^{d_{\text{bg}}}$. This design choice is motivated by the fact that the background is mostly constant between different frames, apart from small variations such as lighting.

3D Gaussian UV Decoder. We use a convolutional decoder to convert the resulting latent variable $\mathbf{Z}$ into a high-resolution UV feature map $\mathbf{U} \in \mathbb{R}^{H \times W \times 23}$. Specifically, for each UV location (u, v) , the output UV features are mapped to the corresponding 3D Gaussian attributes as follows:

$$\boldsymbol{\mu}_{u,v} = \mathbf{V}(u, v) + \mathbf{U}^{(\mu)}(u, v), \quad \|\mathbf{U}^{(\mu)}(u, v)\| \leq \delta_{\text{max}} \quad (3)$$

$$\mathbf{s}_{u,v} = \exp(\mathbf{U}^{(s)}(u, v)) \quad (4)$$

$$\mathbf{q}_{u,v} = \text{normalize}(\mathbf{U}^{(q)}(u, v)) \quad (5)$$

$$\alpha_{u,v} = \sigma(\mathbf{U}^{(\alpha)}(u, v)) \quad (6)$$

$$\{\mathbf{c}_{u,v}^{(\ell,m)}\} = \mathbf{U}^{(c)}(u, v), \quad (7)$$

where $\mathbf{V}(u, v)$ is the corresponding 3D vertex position on the template mesh $\mathbf{V} \in \mathbb{R}^{(H \times W) \times 3}$, and δ_{max} is the position offset bound (empirically set to 200mm).

Similarly, $\mathbf{z}_{bg}$ is decoded into the background 3D Gaussians using a similar but separate decoder (with $\delta_{\max} = 10\text{mm}$). During training, we find it important to perform a warm-up period of 1000 iterations, where opacity and scale attributes are detached from the gradient backpropagation graph.

3.3 Training Objectives

Our overall training loss consists of reconstruction and regularization terms:

$$\mathcal{L}_{\text{total}} = \lambda_{\text{L1}}\mathcal{L}_{\text{L1}} + \lambda_{\text{LPIPS}}\mathcal{L}_{\text{LPIPS}} + \lambda_{\text{adv}}\mathcal{L}_{\text{adv}} + \lambda_{\text{pos}}\mathcal{L}_{\text{pos}} + \lambda_{\text{mask}}\mathcal{L}_{\text{mask}} + \lambda_{\text{TV}}\mathcal{L}_{\text{TV}}, \quad (8)$$

where the λ hyperparameters control the relative influence of each term.

Reconstruction Losses. Let I and I_{gt} denote the rendered and ground-truth composite images (foreground and background) from a randomly sampled view. We optimize an L1 photometric loss $\mathcal{L}_{\text{L1}} = \|I - I_{\text{gt}}\|_1$ and a multi-scale LPIPS perceptual loss $\mathcal{L}_{\text{LPIPS}}$ [25]. To enhance high-frequency details, we employ a Perceptual Discriminator [18, 62] to compute an adversarial loss $\mathcal{L}_{\text{adv}}$. To maintain training stability, $\mathcal{L}_{\text{adv}}$ is activated only after 240k iterations, and the discriminator operates on random 256×256 spatial crops of the input.

Regularization Terms. To guide the geometry during the warm-up phase, we utilize expression-tracked meshes to regularize the Gaussian positions. Specifically, the loss term $\mathcal{L}_{\text{pos}}$ penalizes the distance between $\mu(u, v)$ and the 3D point $\mathbf{V}^e(u, v)$, derived via barycentric interpolation of the Gaussian UV coordinates on a tracked mesh $\mathbf{V}^e$. Concurrently, a silhouette loss $\mathcal{L}_{\text{mask}}$ minimizes the discrepancy between the rendered foreground alpha map and the ground-truth segmentation mask M_{gt} . The weights for both $\mathcal{L}_{\text{pos}}$ and $\mathcal{L}_{\text{mask}}$ are decayed over the course of training. Finally, a Total Variation loss $\mathcal{L}_{\text{TV}}$ is applied to the rendered UV-space colors [32] to encourage spatial smoothness and prevent surface holes.

3.4 Two-Stage Training

Training our model on a large dataset of high-resolution input images can be computationally expensive as the transformer scales quadratically with the number of input tokens. Therefore, we adopt a two-stage training strategy: we first train on $2\times$ downsampled images, followed by a high-resolution finetuning stage that uses the native image resolution. During high-resolution finetuning we make two key modifications:

- **Region-Specific Losses.** We introduce region-specific losses for areas with high-frequency details. Specifically, we extract image crops around the eyes and mouth using the head canonical coordinate frame and supervise with additional multi-scale LPIPS perceptual losses.

- **Multi-Resolution Loss Strategy.** We observed that applying global perceptual and discriminator losses directly at native resolution leads to training instability, likely due to noisy high-frequency gradients. Therefore, to ensure stable convergence, we compute these losses on $2\times$ downsampled outputs.

More implementation details are provided in the Supplementary Material.

4 Experiments

4.1 Datasets

Internal Multi-View Head Dataset. Our internal dataset contains over 10 000 unique participants recorded in a head-focused rig with 16 calibrated RGB cameras. The images have a resolution of 1000×750 . Participants were asked to perform various facial expressions and speech sequences. We train on 10 000 subjects with 100 frames per subject sampled for expression diversity. We evaluate on 100 frames from 50 validation subjects. All participants provided written informed consent for use of their data. The internal dataset will not be publicly released. Images shown in this paper are from subjects who provided explicit written consent for the use of their images in the publication and visualizations.

Ava-256 Dataset. We also finetune and evaluate our model on the public dataset Ava-256 [47], which contains high-resolution multi-view head captures. We use the 4TB version of the dataset, containing 256 subjects recorded with 80 RGB cameras. The images have a resolution of 1024×667 . Following Avat3r’s setup [34], we train on 244 subjects with 1000 frames per subject and evaluate on 12 validation subjects with around 2000 sampled frames in total. Unlike previous methods restricted to frontal views, we sample views from all 80 cameras.

4.2 Experimental Setup

Training. We train our model on the internal dataset for 900K steps with $2\times$ downsampled images (batch size 64) in the first stage, followed by 200K steps of full-resolution finetuning (batch size 32). We use 10 input images for this setup. Training with the Adam optimizer [30] with a learning rate of 2×10^{-4} converges in approximately 10 days on 16 H100 GPUs with `bf16` precision. We perform full-resolution finetuning on Ava-256 for 200K steps with 16 input images, converging in less than a day.

Metrics. We evaluate rendering quality with three image metrics: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS) [82]. Following Avat3r [34], we also report two face-specific metrics: Average Keypoint Distance (AKD) which measures the distance in pixels between 2D keypoints estimated from PIPNet [24], and cosine similarity (CSIM) of ArcFace identity embeddings [12].




Fig. 3. Visual Comparison on Ava-256. Our method produces sharper reconstructions with better identity preservation compared to prior work. Increasing the number of views permits reconstruction of details like earrings, hair and skin texture. Additionally, our background model successfully captures intricate head-boundary details that previous foreground-masking techniques typically discard; we only use the background model during training. Due to GPU memory constraints during training, Avat3r is limited to a maximum of 6 views.

4.3 Baseline

For our baseline comparison, we evaluate our method against Avat3r, representing the state-of-the-art in feed-forward 3D head reconstruction.

Avat3r [34]. Avat3r uses DUST3R [70] position maps and Sapiens features [29] to reconstruct head avatars from sparse views. Since the official code is not released, we carefully reimplemented Avat3r and made some modifications to enable fair comparison to our work. Specifically, we adapt Avat3r from a sparse-view self-reenactment setting to a large-scale 3D head reconstruction pipeline by (1) removing the expression rigging module and (2) providing time-synchronized multi-view images as input. Furthermore, the original method precomputes DUST3R and Sapiens features for only 10 frames per subject, whereas our experimental setup scales this to 1000 frames per subject. As extracting DUST3R position maps at this scale is computationally prohibitive, we instead compute them using the more efficient VGGT [69].

4.4 Baseline Comparison

We compare our method with Avat3r [34] on our internal dataset (Internal10K) and Ava-256 dataset [47] using different metrics. In Ava-256 experiments, both our

Internal10K							
	#V	#G	PSNR↑	SSIM↑	LPIPS↓	AKD↓	CSIM↑
Avat3r	4	0.8M	24.10	0.830	0.371	12.78	0.738
Avat3r	6	1.1M	24.37	0.831	0.362	12.99	0.787
Ours	10	65K	29.25	0.821	0.117	10.39	0.922

Ava-256							
	#V	#G	PSNR↑	SSIM↑	LPIPS↓	AKD↓	CSIM↑
Avat3r	4	0.8M	21.54	0.788	0.370	6.39	0.738
Avat3r	6	1.1M	22.31	0.795	0.359	6.10	0.775
Ours	4	65K	23.64	0.805	0.178	5.09	0.833
Ours	6	65K	24.62	0.815	0.161	4.70	0.873
Ours	16	65K	26.13	0.831	0.110	4.27	0.914

Table 1. Quantitative Comparison. We evaluate both our method and Avat3r [34] on two datasets: our Internal10K dataset, and Ava-256 [47] (finetuned from the Internal10K models), across varying numbers of input views (#V). Our approach significantly outperforms the baseline with major gains in rendering quality (PSNR, LPIPS) and facial fidelity (AKD, CSIM) metrics. Notably, we achieve these improvements while requiring substantially fewer 3D Gaussians (#G). Also, Avat3r is bottlenecked by GPU memory limits at 6 views on our training GPU budget.

model and Avat3r are finetuned from the models pretrained on Internal10K. As shown in Tab. 1, our method significantly outperforms the baseline, with major gains in rendering quality (PSNR, LPIPS) and face fidelity (AKD, CSIM) metrics, while requiring more than an order of magnitude fewer Gaussians. Qualitative results in Fig. 3 also show renders with much higher fidelity and sharper details for our method, capturing fine hair and facial details, with sharper eye and mouth regions than the baseline. Notably, Avat3r is bottle-necked by GPU memory limits at 6 views, whereas our efficient formulation allows for many more inputs. More results and video visualizations are provided in the Supplementary Material.

4.5 Analysis and Ablation Study

We ablate the proposed components of our overall method in Fig. 4. We conduct extensive analysis to validate our design choices and characterize the scaling behavior of our method. We organize our analysis around five key aspects: number of training identities, number of input views, representation capacity, number of target views used for supervision, and our high-resolution finetuning strategy. Moreover, we study the effect of the background model, region-specific perceptual losses, and multi-resolution loss.

Number of training identities. As depicted in Fig. 4a, PSNR improves 1.7–1.8 dB per doubling of subjects up to 2K, with reduced performance gains after 4K. Identity preservation (CSIM) scales more steeply, from 0.441 (250 subjects) to 0.888 (8K). Models trained on <1K subjects fail catastrophically on out-of-distribution faces. Visual comparisons are provided in Fig. 5.

Number of input views. We train our model with an increasing number of random input views. For evaluation, we use a fixed set of views with varying size per experiment, e.g. the frontal view for the monocular case. As shown in Fig. 4b and Fig. 6, more views improve the quality, with diminishing gains after 8 views. Notably, our monocular model is able to convincingly reconstruct the subject.




Fig. 4. Ablation study on a single-stage model trained for 500K steps with 10K subjects, 10 input views, 32×32 latent and 256×256 Gaussian UV resolution unless stated otherwise. **(a) Training data scaling:** log-linear improvement up to 2K subjects, diminishing returns beyond 4K. **(b) Input view scaling:** quality improves with more views, with diminishing returns after 8. **(c) Model capacity:** increasing latent resolution yields larger gains than increasing Gaussian UV resolution. **(d) Template mesh type:** a fixed neutral mesh outperforms expression-tracked meshes. **(e) Number of target views:** more supervision views improve geometric consistency **(f) High-resolution finetuning:** Effect of different components in our second-stage training strategy. **(g) Background model ablation:** Our background modeling permits reconstruction of foreground details such as strands of hair, without background artifacts caused by imperfect image matting techniques.

Latent and Gaussian UV resolution. As shown in Fig. 4c, increasing the latent resolution yields more significant improvements than the UV resolution.

Template mesh type. Our ablation in Fig. 4d shows that using a fixed neutral template significantly outperforms using an expression-tracked one.

Number of target views. Adding more supervision views improves all the metrics by encouraging better view-consistency as shown in Fig. 4e.

High-resolution finetuning. As discussed in Sec. 3.4, after our low-resolution training has converged, we finetune our model on high-resolution target and input images. In Fig. 4f we examine our design choices for the second training stage: *Region-specific losses.* Our eye and mouth crop losses improve reconstruction quality in critical facial features where human perception is most sensitive. *Finetuning perceptual losses.* Downsampling the rendered and real inputs to the Discriminator and LPIPS losses is important for the stability of the adversarial losses. Naively computing LPIPS on high resolution images leads to quality regressions.

Background Gaussian model. As shown in Fig. 4g, removing the background model degrades rendering quality for regions with semi-transparent areas or fine




Fig. 5. Training data scaling. Models trained on fewer subjects fail to generalize to reconstruction of novel identities. At 250 subjects, facial features and hair color deviate significantly from ground truth. The reconstruction quality improves with more training data. On this validation set, the quality improves marginally after 4K subjects.




Fig. 6. Impact of the number of input views. Reconstruction quality scales naturally with the number of input images. A single frontal view ($N = 1$) yields blurry results, identity drift, and fails to recover shirt text. However, adding more views progressively resolves these ambiguities, yielding clear improvements in fine details like the teeth and hair.

foreground elements. Without explicit background modeling, the model hallucinates background artifacts, such as discoloration in the hair. Since computed foreground masks are not pixel-perfect and view-consistent, background details are often combined with foreground elements within masked ground truth images used for supervision. By eliminating the need for these masks, our dedicated background model prevents artifacts from background regions, such as the multi-camera capture rig, and allows the foreground model to focus more Gaussians on facial details.




Fig. 7. Downstream Applications. (a) **Text-driven identity generation:** Novel identities generated by a text-conditioned diffusion model trained on our latents. Using nearest-neighbor face similarity, we verify that these synthesized subjects do not exist in our training set. (b) **Blendshape-driven latent animation:** We train a network conditioned on expression blendshapes, applying a target expression (blue box) to a source identity (green box). The model successfully animates the faces while preserving the subject’s appearance. Both applications operate entirely within the latent space, requiring no per-subject fine-tuning.

4.6 Downstream Applications

Text-driven Identity Generation. Our compact, yet information-rich latent space enables a range of downstream applications, such as novel identity generation. To show this, we train a text-conditioned DiT [52] on a large dataset of latents $\mathbf{Z}$ precomputed from our base model. At inference time, we sample latents and decode them into Gaussians using our frozen decoder. Fig. 7a shows randomly sampled identities from our trained model. Based on face-similarity analysis, we have confirmed these identities do not appear in the training set.

Blendshape-driven Latent Animation. We showcase facial animation controlled by expression blendshapes, operating entirely within our latent space. Given triplets (Z_n, Z_b, b) of a neutral latent Z_n , an expression latent Z_b of the same subject and the corresponding blendshape coefficients b , we train a transformer F_θ to predict the target expression latent $\hat{Z}_b = F_\theta(Z_n, b)$ with supervised losses on the latents, Gaussians and renders. Results are displayed in Fig. 7b.

5 Conclusion

We present *HeadsUp*, a highly scalable feed-forward approach for state-of-the-art reconstruction of 3D Gaussian heads from multi-camera studio captures. By anchoring a compact set of 3D Gaussians to a neutral UV template and employing a lightweight cross-attention transformer, our method achieves photorealistic rendering while gracefully scaling with respect to the number of input views. Crucially, our explicit background modeling eliminates the reliance on imperfect

segmentation masks. Combined with a two-stage training strategy and region-specific perceptual losses, this enables the faithful reconstruction of challenging, high-frequency details such as hair strands and jewelry. We show that *HeadsUp* successfully scales to a production-level dataset of 10 000 unique identities, delivering state-of-the-art reconstruction quality and exhibiting robust generalization to unseen subjects. Finally, beyond feed-forward reconstruction, we show that our learned latent space enables downstream generative applications, including the text-driven synthesis of novel identities and blendshape-driven latent animation.

Additional details on the methodology and experimental setup, along with qualitative results including rendered images and videos, are provided in the supplementary material.

Acknowledgements

We thank Simon Biland, Armin Kappeler and Rick Zhang for their support and valuable feedback.

References

1. An, S., Xu, H., Shi, Y., Song, G., Ogras, U.Y., Luo, L.: PanoHead: Geometry-Aware 3D Full-Head Synthesis in 360°. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 20950–20959 (2023)
2. Athar, S., Xu, Z., Sunkavalli, K., Shechtman, E., Shu, Z.: RigNeRF: Fully Controllable Neural 3D Portraits. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 20364–20373 (2022)
3. Bhattarai, A.R., Nießner, M., Sevastopolsky, A.: TriPlaneNet: An Encoder for EG3D Inversion. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 3533–3542 (2024)
4. Bühler, M.C., Li, G., Wood, E., Helminger, L., Chen, X., Shah, T., Wang, D., Garbin, S., Orts-Escolano, S., Hilliges, O., Lagun, D., Riviere, J., Gotardo, P., Beeler, T., Meka, A., Sarkar, K.: Cafca: High-quality Novel View Synthesis of Expressive Faces from Casual Few-shot Captures. In: SIGGRAPH Asia 2024 Conference Papers. ACM (2024). <https://doi.org/10.1145/3680528.3687580>
5. Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L.J., Tremblay, J., Khamis, S., et al.: Efficient Geometry-aware 3D Generative Adversarial Networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16123–16133 (2022)
6. Charatan, D., Li, S., Tagliasacchi, A., Sitzmann, V.: pixelSplat: 3D Gaussian Splats from Image Pairs for Scalable Generalizable 3D Reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 19457–19467 (2024)
7. Chen, J., Li, C., Zhang, J., Zhu, L., Huang, B., Chen, H., Lee, G.H.: Generalizable Human Gaussians from Single-View Image. arXiv preprint arXiv:2406.06050 (2024)
8. Chen, Y., Xu, H., Zheng, C., Zhuang, B., Pollefeys, M., Geiger, A., Cham, T.J., Cai, J.: MVSplat: Efficient 3D Gaussian Splatting from Sparse Multi-View Images. In: European Conference on Computer Vision (ECCV). pp. 370–386. Springer (2024)

9. Chu, X., Harada, T.: Generalizable and Animatable Gaussian Head Avatar. *NeurIPS* (2024)
10. Chu, X., Li, Y., Zeng, A., Yang, T., Lin, L., Liu, Y., Harada, T.: GPAvatar: Generalizable and Precise Head Avatar from Image(s). *arXiv preprint arXiv:2401.10215* (2024)
11. Debevec, P., Hawkins, T., Tchou, C., Duiker, H.P., Sarokin, W.: Acquiring the Reflectance Field of a Human Face. In: *SIGGRAPH*. New Orleans, LA (Jul 2000), <http://ict.usc.edu/pubs/Acquiring%20the%20Re%EF%AC%82ectance%20Field%20of%20a%20Human%20Face.pdf>
12. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: ArcFace: Additive Angular Margin Loss for Deep Face Recognition. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4690–4699 (2019)
13. Deng, Y., Wang, D., Ren, X., Chen, X., Wang, B.: Portrait4D: Learning One-Shot 4D Head Avatar Synthesis using Synthetic Data. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2024)
14. Deng, Y., Wang, D., Wang, B.: Portrait4D-v2: Pseudo Multi-View Data Creates Better 4D Head Synthesizer. In: *European Conference on Computer Vision (ECCV)*. pp. 303–321. Springer (2024)
15. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *ICLR* (2021)
16. Gafni, G., Thies, J., Zollhöfer, M., Nießner, M.: Dynamic Neural Radiance Fields for Monocular 4D Facial Avatar Reconstruction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 8649–8658 (2021)
17. Giebenhain, S., Kirschstein, T., Georgopoulos, M., Rünz, M., Agapito, L., Nießner, M.: MonoNPHM: Dynamic Head Reconstruction from Monocular Videos. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 21051–21061 (2024)
18. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative Adversarial Networks. In: *Advances in neural information processing systems*. pp. 2672–2680 (2014), <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>
19. Gu, Y., Xu, H., Xie, Y., Song, G., Shi, Y., Di, Y., Ye, P., et al.: DiffPortrait3D: Controllable Diffusion for Zero-Shot Portrait View Synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2024)
20. He, Y., Gu, X., Ye, X., Xu, C., Zhao, Z., Dong, Y., Yuan, W., Dong, Z., Bo, L.: LAM: Large Avatar Model for One-shot Animatable Gaussian Head. In: *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*. pp. 1–13 (2025)
21. Hong, Y., Zhang, K., Gu, J., Bi, S., Zhou, Y., Liu, D., Liu, F., Sunkavalli, K., Bui, T., Tan, H.: LRM: Large Reconstruction Model for Single Image to 3D. *arXiv preprint arXiv:2311.04400* (2023)
22. International Organization for Standardization: ISO 7250-1:2017 Basic human body measurements for technological design — Part 1: Body measurement definitions and landmarks (2017), <https://www.iso.org/standard/65246.html>, accessed: 2026-02-16

23. Ji, X., Weiss, S., Kansy, M., Naruniec, J., Cao, X., Solenthaler, B., Bradley, D.: FastGHA: Generalized Few-Shot 3D Gaussian Head Avatars with Real-Time Animation. arXiv preprint arXiv:2601.13837 (2026)
24. Jin, H., Liao, S., Shao, L.: Pixel-in-Pixel Net: Towards Efficient Facial Landmark Detection in the Wild. *International Journal of Computer Vision* **129**(12), 3174–3194 (2021)
25. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual Losses for Real-Time Style Transfer and Super-Resolution. In: *European Conference on Computer Vision* (2016)
26. Kant, Y., Weber, E., Kim, J.K., Khirodkar, R., Zhaoen, S., Martinez, J., Gilitschenski, I., Saito, S., Bagautdinov, T.: Pippo: High-Resolution Multi-View Humans from a Single Image. In: *CVPR*. pp. 16418–16429 (2025)
27. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM TOG* **42**(4) (July 2023), <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>
28. Khakhulin, T., Sklyarova, V., Lempitsky, V., Zakharov, E.: Realistic One-shot Mesh-based Head Avatars. arXiv preprint arXiv:2206.08343 (2022)
29. Khirodkar, R., Bagautdinov, T., Martinez, J., Zhaoen, S., James, A., Selednik, P., Anderson, S., Saito, S.: Sapiens: Foundation for Human Vision Models. In: *ECCV*. pp. 206–228. Springer (2024)
30. Kingma, D.P.: Adam: A Method for Stochastic Optimization. arXiv preprint arXiv:1412.6980 (2014)
31. Kirillov, A., Mintun, E., Ravi, N., et al.: Segment Anything. arXiv preprint arXiv:2304.02643 (2023)
32. Kirschstein, T., Giebenhain, S., Tang, J., Georgopoulos, M., Nießner, M.: GGHead: Fast and Generalizable 3D Gaussian Heads. In: *SIGGRAPH Asia 2024 Conference Papers*. pp. 1–11 (2024)
33. Kirschstein, T., Qian, S., Giebenhain, S., Walter, T., Nießner, M.: NeRsemble: Multi-view Radiance Field Reconstruction of Human Heads. *ACM TOG* **42**(4), 1–14 (2023)
34. Kirschstein, T., Romero, J., Sevastopolsky, A., Nießner, M., Saito, S.: Avat3r: Large animatable gaussian reconstruction model for high-fidelity 3d head avatars. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 12089–12100 (October 2025)
35. Kwon, Y., Fang, B., Lu, Y., Dong, H., Zhang, C., Vicente Carrasco, F., Mosella-Montoro, A., Xu, J., Takagi, S., Kim, D., Prakash, A., De la Torre, F.: Generalizable Human Gaussians for Sparse View Synthesis. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (2024)
36. Lawrence, J., Goldman, D.B., Achar, S., Blascovich, G.M., Desloge, J.G., Fortes, T., Gomez, E.M., Häberling, S., Hoppe, H., Huibers, A., Knaus, C., Kuschak, B., Martin-Brualla, R., Nover, H., Russell, A.I., Seitz, S.M., Tong, K.: Project Starline: A High-Fidelity Telepresence System. *ACM Transactions on Graphics (Proc. of SIGGRAPH Asia)* **40**(6) (2021)
37. Li, H., Chen, C., Shi, T., Qiu, Y., An, S., Chen, G., et al.: SphereHead: Stable 3D Full-head Synthesis with Spherical Tri-plane Representation. In: *European Conference on Computer Vision (ECCV)*. Springer (2024)
38. Li, J., Cao, C., Schwartz, G., Khirodkar, R., Saito, S., et al.: URAvatar: Universal Relightable Gaussian Codec Avatars. In: *SIGGRAPH Asia 2024 Conference Papers*. ACM (2024)
39. Li, P., He, Y., Hu, Y., Dong, Y., Yuan, W., Liu, Y., Zhu, S., Cheng, G., Dong, Z., Guo, Y.: PanoLAM: Large Avatar Model for Gaussian Full-Head Synthesis from One-shot Unposed Image. arXiv preprint arXiv:2509.07552 (2025)

40. Li, P., Zheng, W., Liu, Y., Yu, T., Li, Y., Qi, X., Chi, X., Xia, S., Cao, Y.P., Xue, W., et al.: PSHuman: Photorealistic Single-Image 3D Human Reconstruction Using Cross-Scale Multiview Diffusion and Explicit Remeshing. In: Proceedings of the computer vision and pattern recognition conference. pp. 16008–16018 (2025)
41. Li, T., Bolkart, T., Black, M.J., Li, H., Romero, J.: Learning a Model of Facial Shape and Expression from 4D Scans. *ACM TOG* **36**(6), 194–1 (2017)
42. Liang, H., Ge, Z., Tiwari, A., Majee, S., Godaliyadda, G.M.D., Veeraraghavan, A., Balakrishnan, G.: FastAvatar: Instant 3D Gaussian Splatting for Faces from Single Unconstrained Poses. *arXiv preprint arXiv:2508.18389* (2025)
43. Lombardi, S., Saragih, J., Simon, T., Sheikh, Y.: Deep Appearance Models for Face Rendering. *ACM Transactions on Graphics (TOG)* **37**(4), 1–13 (2018)
44. Lu, Y., Dong, J., Kwon, Y., Zhao, Q., Dai, B., De la Torre, F.: GAS: Generative Avatar Synthesis from a Single Image. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12883–12893 (2025)
45. Lyu, W., Zhou, Y., Yang, M.H., Shu, Z.: FaceLift: Learning Generalizable Single Image 3D Face Reconstruction from Synthetic Heads. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12691–12701 (2025)
46. Ma, S., Simon, T., Saragih, J., Wang, D., Li, Y., De La Torre, F., Sheikh, Y.: Pixel Codec Avatars. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 64–73 (2021)
47. Martinez, J., Kim, E., Romero, J., Bagautdinov, T., Saito, S., Yu, S.I., Anderson, S., Zollhöfer, M., Wang, T.L., Bai, S., Li, C., Wei, S.E., Joshi, R., Borsos, W., Simon, T., Saragih, J., Theodosis, P., Greene, A., Josyula, A., Maeta, S.M., Jewett, A.I., Venshtain, S., Heilman, C., Chen, Y.T., Fu, S., Elshaer, M.E.A., Du, T., Wu, L., Chen, S.C., Kang, K., Wu, M., Emad, Y., Longay, S., Brewer, A., Shah, H., Booth, J., Koska, T., Haidle, K., Andromalos, M., Hsu, J., Dauer, T., Selednik, P., Godisart, T., Ardisson, S., Cipperly, M., Humberston, B., Farr, L., Hansen, B., Guo, P., Braun, D., Krenn, S., Wen, H., Evans, L., Fadeeva, N., Stewart, M., Schwartz, G., Gupta, D., Moon, G., Guo, K., Dong, Y., Xu, Y., Shiratori, T., Prada, F., Pires, B.R., Peng, B., Buffalini, J., Trimble, A., McPhail, K., Schoeller, M., Sheikh, Y.: Codec Avatar Studio: Paired Human Captures for Complete, Driveable, and Generalizable Avatars. *NeurIPS* (2024)
48. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In: European Conference on Computer Vision (ECCV) (2020)
49. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: DINOv2: Learning Robust Visual Features without Supervision. *arXiv preprint arXiv:2304.07193* (2023)
50. Oroz, A., Nießner, M., Kirschstein, T.: PercHead: Perceptual Head Model for Single-Image 3D Head Reconstruction & Editing. *arXiv preprint arXiv:2511.02777* (2025)
51. Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable Neural Radiance Fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 5865–5874 (2021)
52. Peebles, W., Xie, S.: Scalable Diffusion Models with Transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 4195–4205 (2023)

53. Prinzler, M., Gotardo, P., Tang, S., Bolkart, T.: Feed-forward gaussian registration for head avatar creation and editing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2026)
54. Qian, S., Kirschstein, T., Schoneveld, L., Davoli, D., Giebenhain, S., Nießner, M.: GaussianAvatars: Photorealistic Head Avatars with Rigged 3D Gaussians. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
55. Qiu, L., Gu, X., Li, P., Zuo, Q., Shen, W., Zhang, J., Qiu, K., Yuan, W., Chen, G., Dong, Z., Bo, L.: LHM: Large Animatable Human Reconstruction Model from a Single Image in Seconds. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2025)
56. Qiu, L., Li, P., Zuo, Q., Gu, X., Dong, Y., Yuan, W., Zhu, S., Han, X., Chen, G., Dong, Z.: PF-LHM: 3D Animatable Avatar Reconstruction from Pose-free Articulated Human Images. arXiv preprint arXiv:2506.13766 (2025)
57. Roich, D., Mokady, R., Bermano, A.H., Cohen-Or, D.: Pivotal Tuning for Latent-based Editing of Real Images. ACM Transactions on Graphics (TOG) **42**(1), 1–13 (2022)
58. Saito, S., Schwartz, G., Simon, T., Li, J., Nam, G.: Relightable Gaussian Codec Avatars. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 130–141 (2024)
59. Shao, Z., Wang, Z., Li, Z., Wang, D., Lin, X., Zhang, Y., Fan, M., Wang, Z.: SplattingAvatar: Realistic Real-Time Human Avatars with Mesh-Embedded Gaussian Splatting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16922–16932 (2024)
60. Sitzmann, V., Rezchikov, S., Freeman, W.T., Tenenbaum, J.B., Durand, F.: Light Field Networks: Neural Scene Representations with Single-Evaluation Rendering. In: Proc. NeurIPS (2021)
61. Skorokhodov, I., Siarohin, A., Xu, Y., Ren, J., Lee, H.Y., Wonka, P., Tulyakov, S.: 3D Generation on ImageNet. arXiv preprint arXiv:2303.01416 (2023)
62. Sungatullina, D., Zakharov, E., Ulyanov, D., Lempitsky, V.: Image Manipulation with Perceptual Discriminators. In: Proceedings of the European Conference on Computer Vision (ECCV) (September 2018)
63. Szymanowicz, S., Rupprecht, C., Vedaldi, A.: Splatter Image: Ultra-Fast Single-View 3D Reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10208–10217 (2024)
64. Tang, J., Davoli, D., Kirschstein, T., Schoneveld, L., Niessner, M.: GAF: Gaussian Avatar Reconstruction from Monocular Videos via Multi-view Diffusion. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 5546–5558 (2025)
65. Teotia, K., Kim, H., Garrido, P., Habermann, M., Elgharib, M., Theobalt, C.: GaussianHeads: End-to-End Learning of Drivable Gaussian Head Avatars from Coarse-to-fine Representations. ACM Transactions on Graphics (SIGGRAPH Asia) **43**(6) (2024)
66. Teotia, K., Rhodin, H., Mendiratta, M., Kim, H., Habermann, M., Theobalt, C.: Audio Driven Universal Gaussian Head Avatars. In: Proceedings of the SIGGRAPH Asia 2025 Conference Papers. pp. 1–12 (2025)
67. Tran, P., Zakharov, E., Ho, L.N., Tran, A.T., Hu, L., Li, H.: VODOO 3D: Volumetric Portrait Disentanglement for One-Shot 3D Head Reenactment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)

68. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention Is All You Need. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 30. Curran Associates, Inc. (2017)
69. Wang, J., Chen, M., Karaev, N., Vedaldi, A., Rupprecht, C., Novotny, D.: VGGT: Visual Geometry Grounded Transformer. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2025)
70. Wang, S., Leroy, V., Cabon, Y., Chidlovskii, B., Revaud, J.: DUS3R: Geometric 3D Vision Made Easy. In: *CVPR*. pp. 20697–20709 (2024)
71. Xiang, J., Gao, X., Guo, Y., Zhang, J.: FlashAvatar: High-fidelity head avatar with efficient gaussian embedding. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1802–1812 (2024)
72. Xu, Y., Shi, Z., Yifan, W., Chen, H., Yang, C., Peng, S., Shen, Y., Wetzstein, G.: GRM: Large Gaussian Reconstruction Model for Efficient 3D Reconstruction and Generation. In: *European Conference on Computer Vision*. pp. 1–20. Springer (2024)
73. Xue, Y., Xie, X., Marin, R., Pons-Moll, G.: Human-3Diffusion: Realistic Avatar Creation via Explicit 3D Consistent Diffusion Models. *Advances in Neural Information Processing Systems* **37**, 99601–99645 (2024)
74. Yang, J., Wu, T., Fogarty, K., Zhong, F., Oztireli, C.: PSHead: 3D Head Reconstruction from a Single Image with Diffusion Prior and Self-Enhancement. In: *Computer Graphics Forum*. p. e70279. Wiley Online Library (2025)
75. Yang, Y., Liu, F., Lu, Y., Zhao, Q., Wu, P., Zhai, W., Yi, R., Cao, Y., Ma, L., Zha, Z.J., et al.: SIGMAN: Scaling 3D Human Gaussian Generation with Millions of Assets. *ICCV* (2025)
76. Ye, Z., Zhong, T., Ren, Y., Yang, J., Li, W., Huang, J., Jiang, Z., He, J., Huang, R., Liu, J., Zhang, C., Yin, X., Ma, Z., Zhao, Z.: Real3D-Portrait: One-shot Realistic 3D Talking Portrait Synthesis. In: *International Conference on Learning Representations (ICLR)* (2024), https://openreview.net/forum?id=7ERQPyR2eb_spotlight
77. Yuan, Z., Zhu, Y., Li, Y., Liu, H., Yuan, C.: Make Encoder Great Again in 3D GAN Inversion through Geometry and Occlusion-Aware Encoding. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 2437–2447 (2023)
78. Zhang, B., Cheng, Y., Wang, C., Zhang, T., Yang, J., Tang, Y., Zhao, F., Chen, D., Guo, B.: RodinHD: High-Fidelity 3D Avatar Generation with Diffusion Models. In: *European Conference on Computer Vision*. pp. 465–483. Springer (2024)
79. Zhang, J., Gao, Y., Zhan, J., Wang, W., Zhang, Y., Zhao, H., Zhang, L.: High-Quality 3D Head Reconstruction from Any Single Portrait Image. *arXiv preprint arXiv:2503.08516* (2025)
80. Zhang, J., Li, X., Zhang, Q., Cao, Y., Shan, Y., Liao, J.: HumanRef: Single Image to 3D Human Generation via Reference-Guided Diffusion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1844–1854 (2024)
81. Zhang, K., Bi, S., Tan, H., Xiangli, Y., Zhao, N., Sunkavalli, K., Xu, Z.: GS-LRM: Large Reconstruction Model for 3D Gaussian Splatting. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (2024)
82. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In: *CVPR*. pp. 586–595 (2018)

83. Zhao, X., Sun, J., Wang, L., Suo, J., Liu, Y.: InvertAvatar: Incremental GAN Inversion for Generalized Head Avatars. In: ACM SIGGRAPH 2024 Conference Papers. pp. 1–10 (2024)
84. Zhao, Z., Bao, Z., Li, Q., Qiu, G., Liu, K.: PSAvatar: A Point-based Shape Model for Real-Time Head Avatar Animation with 3D Gaussian Splatting. arXiv preprint arXiv:2401.12900 (2024)
85. Zheng, S., Tu, H., Zhou, B., Shao, R., Liu, B., Zhang, S., Nie, L., Liu, Y.: GPS-Gaussian: Generalizable Pixel-wise 3D Gaussian Splatting for Real-time Human Novel View Synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
86. Zheng, X., Wen, C., Li, Z., Zhang, W., Su, Z., Chang, X., Zhao, Y., Lv, Z., Zhang, X., Zhang, Y., et al.: HeadGAP: Few-shot 3D head avatar via generalizable gaussian priors. In: 2025 International Conference on 3D Vision (3DV). pp. 946–957. IEEE (2025)
87. Zheng, Y., Abrevaya, V.F., Bühler, M.C., Chen, X., Black, M.J., Hilliges, O.: I M Avatar: Implicit Morphable Head Avatars from Videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 13545–13555 (2022)
88. Zheng, Y., Yifan, W., Wetzstein, G., Black, M.J., Hilliges, O.: PointAvatar: Deformable Point-based Head Avatars from Videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 21057–21067 (2023)
89. Zielonka, W., Bolkart, T., Thies, J.: Instant volumetric head avatars. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4574–4584 (2023)