

# ScaffoldAvatar: High-Fidelity Gaussian Avatars with Patch Expressions

SHIVANGI ANEJA, Technical University of Munich, Germany and DisneyResearch|Studios, Switzerland  
SEBASTIAN WEISS, DisneyResearch|Studios, Switzerland  
IRENE BAEZA, DisneyResearch|Studios, Switzerland  
PRASHANTH CHANDRAN, DisneyResearch|Studios, Switzerland  
GASPARD ZOSS, DisneyResearch|Studios, Switzerland  
MATTHIAS NIESSNER, Technical University of Munich, Germany  
DEREK BRADLEY, DisneyResearch|Studios, Switzerland

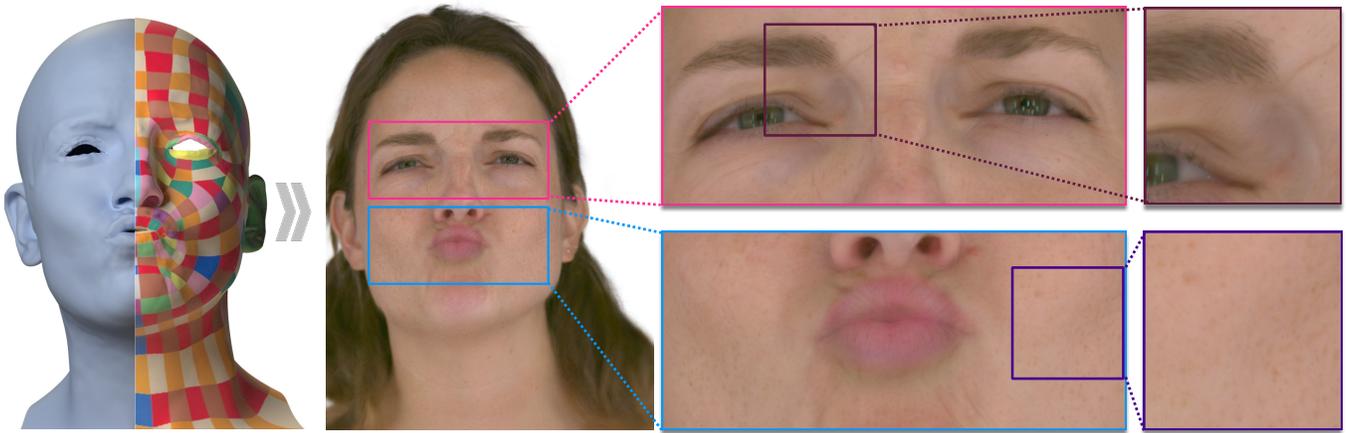


Fig. 1. Given an input mesh with patch-level localized expressions (left), ScaffoldAvatar can synthesize ultra-high fidelity multi-view consistent photorealistic avatars. Our avatars can synthesize high frequency facial details as shown in two-level zoom-ins (right). Our method is capable of generating realistic and high-quality animations including freckles and other fine facial details with real-time rendering.

Generating high-fidelity real-time animated sequences of photorealistic 3D head avatars is important for many graphics applications, including immersive telepresence and movies. This is a challenging problem particularly when rendering digital avatar close-ups for showing character's facial microfeatures and expressions. To capture the expressive, detailed nature of human heads, including skin furrowing and finer-scale facial movements, we propose to couple locally-defined facial expressions with 3D Gaussian splatting to enable creating ultra-high fidelity, expressive and photorealistic

Authors' addresses: Shivangi Aneja, Technical University of Munich, Munich, Germany, shivangi.aneja@tum.de and DisneyResearch|Studios, Zurich, CA, Switzerland, shivangi.aneja@disneyresearch.com; Sebastian Weiss, DisneyResearch|Studios, Zurich, CA, Switzerland, sebastian.weiss@disneyresearch.com; Irene Baeza, DisneyResearch|Studios, Zurich, CA, Switzerland, irene.baeza@disneyresearch.com; Prashanth Chandran, DisneyResearch|Studios, Zurich, CA, Switzerland, prashanthchandran.pc@gmail.com; Gaspard Zoss, DisneyResearch|Studios, Zurich, CA, Switzerland, gaspard.zoss@gmail.com; Matthias Niessner, Technical University of Munich, Germany, niessner@tum.de; Derek Bradley, DisneyResearch|Studios, Zurich, CA, Switzerland, derek.bradley@disneyresearch.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

SIGGRAPH Conference Papers '25, August 10–14, 2025, Vancouver, BC, Canada  
© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 979-8-4007-1540-2/2025/08  
<https://doi.org/10.1145/3721238.3730729>

3D head avatars. In contrast to previous works that operate on a global expression space, we condition our avatar's dynamics on patch-based local expression features and synthesize 3D Gaussians at a patch level. In particular, we leverage a patch-based geometric 3D face model to extract patch expressions and learn how to translate these into local dynamic skin appearance and motion by coupling the patches with anchor points of Scaffold-GS, a recent hierarchical scene representation. These anchors are then used to synthesize 3D Gaussians on-the-fly, conditioned by patch-expressions and viewing direction. We employ color-based densification and progressive training to obtain high-quality results and faster convergence for high resolution 3K training images. By leveraging patch-level expressions, ScaffoldAvatar consistently achieves state-of-the-art performance with visually natural motion, while encompassing diverse facial expressions and styles in real time.

CCS Concepts: • **Computing methodologies** → **Rasterization**; *Procedural animation*.

Additional Key Words and Phrases: Photorealistic 3D avatars, 3D Gaussian Splatting, 3D Face Modelling

## ACM Reference Format:

Shivangi Aneja, Sebastian Weiss, Irene Baeza, Prashanth Chandran, Gaspard Zoss, Matthias Niessner, and Derek Bradley. 2025. ScaffoldAvatar: High-Fidelity Gaussian Avatars with Patch Expressions. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '25)*, August 10–14, 2025, Vancouver, BC, Canada. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3721238.3730729>

## 1 INTRODUCTION

The creation of realistic digital human avatars is a fast-evolving field with many applications in virtual telepresence, movies and entertainment. The resulting digital avatar should be photorealistic, animatable, allow novel viewpoint rendering in real-time, and must preserve a high degree of spatial fidelity, such as fine-scale expression-specific skin details during animation. Developing a 3D representation that can capture high-frequency appearance and accurate motion dynamics of human heads represents a major challenge, in particular under novel close-up views of the face where the skin exhibits many details. We require a representation that can jointly reconstruct the facial motion while simultaneously retaining capacity to facilitate high-fidelity re-rendering.

Recent advances in computer vision and graphics, in particular 3D Gaussian Splatting (3DGS) [Kerbl et al. 2023], has quickly gained popularity in many application domains due to its simplicity and rendering speed, combined with its differentiable formulation that can be optimized to reconstruct a scene given a set of images from different viewpoints. Several works have extended 3DGS to adapt it for multi-view consistent animatable human avatars [Giebenhain et al. 2024b; Qian et al. 2024; Teotia et al. 2024; Xu et al. 2023a], achieving a new state-of-the-art in realistic digital humans. These existing methods achieve most of the criteria required: photorealism, animatable, novel viewpoints, and real-time, but they fall short on actor fidelity. For example, the recent GaussianAvatars [Qian et al. 2024] method opened the door to real-time photoreal human heads, but it did not focus on high actor fidelity, so re-rendering appears blurry when zoomed in and expression-dependent appearance details like wrinkles are not supported. Neural Parametric Gaussian Avatars (NPGA) [Giebenhain et al. 2024b] achieves dynamic expression wrinkles, but still loses fidelity and appear blurry when zoomed in. Gaussian Head Avatars (GHA) [Xu et al. 2023a] make a first attempt at high-fidelity by training on 2K images, which is nearly double the resolution of previous methods. However, as we will show in the results, zoom-in regions still lack the necessary details to recover the actor’s appearance with sufficient fidelity. This lack of fine-scale facial expressions is due to the use of global expression conditioning from existing 3D face models [Gerig et al. 2017; Giebenhain et al. 2024a]. Such global expression spaces lack expressiveness for fine-grained facial regions, and the inaccurate motion representation due to the low-dimension expression codes leads to blurred out appearance. In summary, existing methods lack the representative capacity to re-render photoreal fine scale details in dynamic human faces, which is the problem we tackle in this work.

To this end, we present ScaffoldAvatar, a novel patch-based method for real-time digital avatars that explicitly tackles high spatial fidelity with a novel formulation to increase representative capacity, beyond what is achievable with current state-of-the-art methods. We achieve this by operating on localized patch expressions. Our approach has similarities to Scaffold-GS [Lu et al. 2024], a spatially hierarchical 3DGS approach that gives us local control of small facial regions. Our unique architecture provides local network capacity dedicated to recovering Gaussian parameters for the local regions,

conditioned on learned local features and local expression parameters. The scaffold anchors are attached to patch centers of a tracked 3D mesh, where we use a localized patch-level facial expression, following Chandran *et al.* [2022] who showed that a patch blend-shape model is both very expressive and cheap to compute. The system is trained end-to-end on multi-view dynamic performance sequences with 3K high resolution images. Furthermore, unlike existing methods that use positional gradients for densification, we use view-space color gradients to recover sharper skin texture and employ progressive training for faster convergence.

ScaffoldAvatar has the capacity to represent fine-scale skin details, including dynamic wrinkles, freckles, and high frequency facial details (as depicted in Fig. 1). Since our MLPs are conditioned on patch-based expression blendweights, we can regress expression-dependent appearance like wrinkles and blood flow. Once trained, our method supports real-time inference with control over expression and viewpoint like existing Gaussian-based avatars, but at higher spatial fidelity than current methods. To summarize, we present the following contributions:

- A novel method capable of synthesizing ultra-high fidelity photorealistic avatars, due to our unique network formulation similar to Scaffold-GS with conditioning on localized patch-level facial expressions, which can reconstruct finegrained facial details like freckles and wrinkles.
- By leveraging progressive training and view-space color gradient for densification, our method can synthesize sharp skin texture and converge faster for 3K high resolution images.

## 2 RELATED WORK

We first discuss 3D face and head models that primarily focus on geometry, and then cover more recent photorealistic 3D human head animation techniques.

### 2.1 Geometric 3D Face and Head Models

The seminal work of Blanz and Vetter [1999] introduced a model-based approach to represent variations in human faces using PCA. Thereafter, more advanced face models were introduced, including multilinear models for identity and expression [Bolkart and Wuhler 2015; Brunton et al. 2014], as well as recent methods combining linear shape spaces with articulated head parts [Li et al. 2017]. In contrast to these model-based approaches relying on global expression blendweights, techniques based on local patch blendweights [Chandran et al. 2022; Neumann et al. 2013] can produce much more accurate deformations. For instance, Neumann *et al.* [2013] demonstrate that by separating facial performances into localized deformation components, they can produce high-detail deformations like muscle bulges. Similarly, Chandran *et al.* [2022] show that operating on small patches on the mesh surface result in more accurately retargeted facial performances. Similar to these works, we leverage local patch blendweights to achieve high frequency details with our patch-based 3DGS avatar.

## 2.2 Photorealistic 3D Human Head Animation

Early approaches for reconstructing photorealistic head avatars were based on NeRFs [Mildenhall et al. 2021], which store the radiance field of the scene in a neural network and render novel views of the scene via volume rendering. Recently, 3D Gaussian Splatting (3DGS) [Kerbl et al. 2023] has emerged as an efficient representation for reconstructing fine geometric structures with real-time rendering, leading to a lot of interest in point-based representations for animatable avatars. We discuss both of these approaches below in more detail, in the context of animatable human head avatars.

**2.2.1 NeRF Based Reconstruction.** Gafni et al. [2021] introduced one of the first methods conditioning NeRFs on the expression features from monocular videos. Following this, INSTA [Zielonka et al. 2022] deforms query points to a canonical space by finding the nearest triangle on a FLAME mesh [Li et al. 2017], combined with InstantNGP [Müller et al. 2022] to achieve fast rendering. NeRF-BlendShape [Gao et al. 2022] models a dynamic scene by blending hash tables with 3DMM parameters. AvatarMAV [Xu et al. 2023b] decouples motion and appearance, blending voxel grids only for the motion field. Nersemble [Kirschstein et al. 2023] uses spatio-temporal NeRFs and learns deformation fields for coarse motion and hash grid encodings for fine scale deformations and is trained using high quality multi-view video data.

**2.2.2 3DGS Based Reconstruction.** Zheng et al. [2023] explored point-based representations with differential point splatting by defining a point set in canonical space and learning a deformation field conditioned on FLAME’s expression space for animating the avatar. GaussianAvatars [Qian et al. 2024] is the seminal work proposing a method for dynamic 3D representation of human heads based on 3DGS by rigging the anisotropic 3D Gaussians to the faces of the FLAME mesh. GaussianHeads [Teotia et al. 2024] predict RGB color and opacity in the UV space of FLAME mesh and deform the canonical gaussians using MLP. GaussianHeadAvatars (GHA) [Xu et al. 2023a] optimizes neural 3D Gaussians with MLP-based deformations driven by the global expression space of BFM [Gerig et al. 2017] for handling the dynamics. Very recently, Neural Parametric Gaussian Avatars (NPGA) [Giebenhain et al. 2024b] also learns canonical 3D Gaussians and deforms them using a MLP conditioned on the rich expression space of neural parametric head models [Giebenhain et al. 2024a]. While both GHA and NPGA are capable of producing quality animations with expression-dependent deformations and refining the results with a screen-space CNN, they cannot produce fine-scale facial details when zoomed-into different regions of the avatar’s face.

In contrast to these techniques based on global expression features, we leverage local patch level expression blendweights and predict patch level deformations and Gaussian parameters, thus our method is capable of producing finer details for facial zoom-ins.

## 3 PRELIMINARIES

3DGS [Kerbl et al. 2023] defines the scene as a collection of Gaussian primitives, each parameterized by position  $\mu$ , scale  $s$ , rotation quaternion  $\mathbf{q}$ , opacity  $\alpha$  and color  $\mathbf{c}$ . The scene is then rendered into images with a differentiable rasterizer. This representation, however,

is very expensive in terms of memory to store thousands to millions of primitives. More details can be found in the supplemental material.

To improve scene coverage and avoid redundant Gaussian primitives and redundant parameters, Scaffold-GS [Lu et al. 2024] introduces a hierarchy of primitives. First, the initial scene is discretized into a sparse voxel grid. The center of each voxel  $v$  at position  $\mathbf{x}_v$  is called an *anchor* and is equipped with a local feature vector  $f_v \in \mathbb{R}^F$ . Each visible anchor at point  $\mathbf{x}_v$  spawns  $L$  Gaussian primitives with each position  $\mu_l$  defined as:

$$\{\mu_1, \dots, \mu_L\} = \mathbf{x}_v + \{\mathcal{O}_1, \dots, \mathcal{O}_L\} \cdot s_v, \quad (1)$$

where  $\{\mathcal{O}_1, \dots, \mathcal{O}_L\} \in \mathbb{R}^{L \times 3}$  represents the learnable offsets and  $s_v$  is the scaling factor from the associated anchor. These offsets are learned per anchor. The remaining 3DGS attributes including opacity  $\alpha$ , color  $\mathbf{c}$ , rotation  $\mathbf{q}$ , and scale  $\mathbf{s}$  are predicted by individual but global MLPs  $F_\alpha, F_c, F_q, F_s$  conditioned on the learned anchor feature vector  $f_v$  and view direction  $d_v$ . Note that the parameters of each Gaussian primitive per anchor are predicted at the same time, for example opacity is predicted as:

$$\{\alpha_1, \dots, \alpha_L\} = F_\alpha(f_v, d_v). \quad (2)$$

Colors  $\{c_l\}$ , rotation quaternions  $\{q_l\}$  and scales  $\{s_l\}$  are predicted in a similar fashion. These spawned Gaussian primitives are then rasterized following standard 3DGS [Kerbl et al. 2023].

## 4 DATA CAPTURE AND PREPARATION

We begin by describing the data we captured for each actor, and the required pre-processing steps.

### 4.1 Capture Setup and Mesh Tracking

Our main capture setup consists of 9 synchronized video cameras spread out in the hemisphere in front of the actor (8 of which are used for training, one center camera for validation). For the sake of hardware availability, we use a combination of 12MP and 20MP color cameras, with a combination of 60mm and 85mm lenses, all shooting at 24fps. We capture between 5-8 performance sequences for each actor, where the sequences contain various localized muscle movements, overall facial workouts, natural expressions, and dialog performances. The multi-view sequences are reconstructed to obtain a topology-consistent tracked mesh of the facial skin and head region using the highly accurate tracking method of Wu et al. [2016], extended to multi-camera inputs.

A second capture setup, which consists of 8 static DSLR cameras arranged in 4 stereo pairs, is used to scan 20 blendshapes of each actor [Beeler et al. 2010], which are used to build the tracking model used on the dynamic performances above, as well as to build a local patch-based model that we describe next.

### 4.2 Patch Blendweight Optimization

We leverage the expressive patch-based facial blendshape model of Chandran et al. [2022] to represent dynamic local facial expressions. To this end, the tracked mesh topology is divided into  $P = 432$  small overlapping patches as shown in Fig. 2. For each patch  $p$  in the set of all patches  $\mathcal{P}$ , we build a local blendshape model from the  $K = 20$

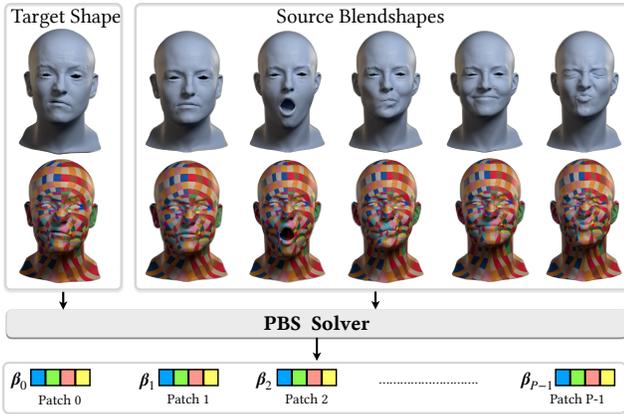


Fig. 2. Patch Blendweight Optimization: Given a target shape from our tracked meshes, we fit patch blendshape model (Eq. 3) and use our PBS Solver to find optimal per-patch blendweights  $\{\beta_i \mid i = 0, 1, \dots, P-1\}$ , where  $P$  is the total number of patches in our mesh. This process is repeated for every frame in the dataset. The patch layout rendering in the bottom row shows that same patch layout is used for all the shapes.

static scans  $\mathcal{S} = \{S_i \mid i = 0, 1, \dots, K-1\}$  as:

$$\mathbf{X}_p^S = p^{S_0} + \sum_{i=1}^{K-1} \beta_{p,i} (p^{S_i} - p^{S_0}), \quad (3)$$

where  $\beta_{p,i} \in \mathbb{R}^{K-1}$  are patch blendweights used to linearly blend the patch blendshapes. Since each patch has its own set of patch blendweights  $\beta_{p,i}$ , this formulation provides more degrees of freedom for representing the expressions than global blendshape models, *i.e.*,  $P \cdot (K-1) = 8208$  total expression parameters vs. 100 of FLAME [Li et al. 2017].

Next we fit the local blendshape models to all the dynamic captured frames for the corresponding actor. Specifically, given a captured shape  $\mathbf{X}^t$  at time  $t = 1 \dots T$ , we solve for all the patch blendweights

$$\beta = \{\beta_{p,i}, p = 1 \dots P, i = 1 \dots K-1\}, \quad (4)$$

such that the resulting deformed patches accurately describe the local skin deformations exhibited in  $\mathbf{X}^t$ , accomplished with a least-squares optimization as:

$$\mathcal{E}_{ls} = \sum_{p \in \mathcal{P}} \left\| \mathbf{X}_p^t - R \mathbf{X}_p^S \right\|_2^2, \quad (5)$$

where  $\mathbf{X}_p^S$  is defined in Eq. 3 and  $R$  is global rigid transformation. The patch coefficients are additionally regularized to remain close to zeros as:

$$\mathcal{E}_{reg} = \sum_{p \in \mathcal{P}} \sum_{i=1}^{K-1} \|\beta_{p,i}\|_2, \quad (6)$$

and to stay consistent across adjacent patches,

$$\mathcal{E}_o = \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{N}(p)} \sum_{i=1}^{K-1} \|\beta_{p,i} - \beta_{q,i}\|_1, \quad (7)$$

where  $\mathcal{N}(p)$  defines the patches neighboring  $p$ . The overall loss for fitting the patch blendshape model to a target shape  $\mathbf{X}^t$  is the

weighted sum of these losses as

$$\mathcal{E}_{PBS} = \lambda_{ls} \mathcal{E}_{ls} + \lambda_{reg} \mathcal{E}_{reg} + \lambda_o \mathcal{E}_o, \quad (8)$$

and final patch blendweights are obtained as:

$$\beta^* = \arg \min_{\beta} \mathcal{E}_{PBS}. \quad (9)$$

This process is repeated for each of the  $T$  shapes in the dynamic data, yielding a sequence of per-frame per-patch expression coefficients  $\{\beta^1, \beta^2, \dots, \beta^T\}$ . Each actor is processed and trained separately.

## 5 SCAFFOLD AVATAR

We propose a novel avatar representation that combines the expressive power of local deformation models for human faces with efficient hierarchical scene rendering (similar to Scaffold-GS). Using the local patch blendshape representation (Sec. 4.2), we assign an anchor to each patch, which spawns multiple Gaussian primitives. These primitives are represented with attributes predicted by MLPs, conditioned on learned anchor features as well as local expression deformation (see Fig. 3 for an overview).

A key contribution of our method is how we build the connection between anchors and the patches of our tracked mesh (Sec. 5.1 and Sec. 5.2) and how to spawn Gaussian primitives around the anchors (Sec. 5.3). Additional optimization and training details are given in Sec. 5.4 and Sec. 5.5, respectively. Hyperparameters are given in the supplementary document.

### 5.1 Defining Patch Centers and Attributes

To determine the center of each patch for a given tracked mesh with  $P$  patches, we first calculate the mean position  $\mathbf{c}_p$  of all vertices belonging to that particular patch  $p$ . Then we designate the closest mesh vertex in Euclidean distance to  $\mathbf{c}_p$  as the patch center  $\mathbf{v}_p$ . The patch centers  $\{\mathbf{v}_p \mid p = 1, \dots, P\}$  are determined once and remain fixed throughout the experiments.

To determine the orientation and position of each patch center  $\mathbf{v}_p$  in global space, we define the TBNP (Tangent-Bitangent-Normal-Position) matrix  $\mathcal{T}_p$  as:

$$\mathcal{T}_p = \begin{bmatrix} \mathbf{t}_p^x & \mathbf{b}_p^x & \mathbf{n}_p^x & \mathbf{v}_p^x \\ \mathbf{t}_p^y & \mathbf{b}_p^y & \mathbf{n}_p^y & \mathbf{v}_p^y \\ \mathbf{t}_p^z & \mathbf{b}_p^z & \mathbf{n}_p^z & \mathbf{v}_p^z \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (10)$$

where  $\mathbf{t}_p$ ,  $\mathbf{b}_p$ , and  $\mathbf{n}_p$  represents the tangent, bitangent and normal vectors at the patch center, and  $\mathbf{v}_p$  denotes the position of patch center in global space. These matrices are defined for every patch, computed for every frame, and later used to displace the anchors and Gaussian primitives.

### 5.2 Rigging Scaffold Anchors to Patches

For a given patch  $p$  we associate a set of anchors in its neighborhood and let these anchors move as the patch moves for different expressions. More specifically, the anchors are static in the local space of its parent patch but change in the global space as the patch moves by transforming them by their parent's patch TBNP matrix.

Let  $\mathcal{A}_p = \{\mathbf{A}_{p,1}, \mathbf{A}_{p,2}, \dots, \mathbf{A}_{p,N}\}$  denote the set of  $N$  anchors associated with the patch center  $\mathbf{v}_p$ , where each anchor contains

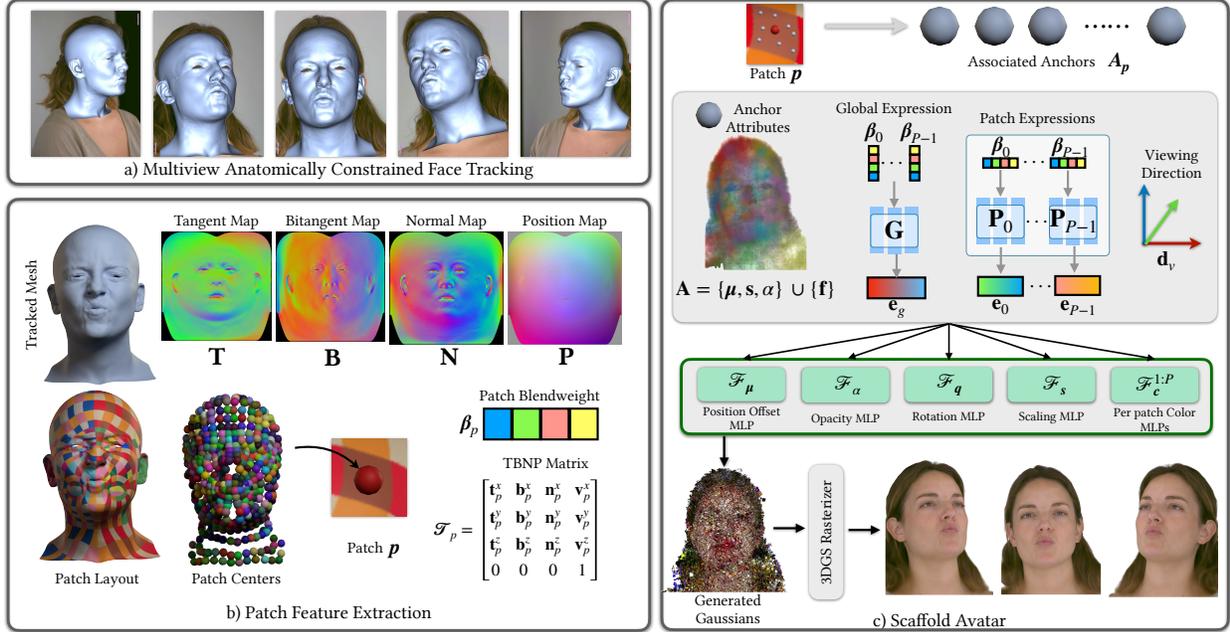


Fig. 3. **Method Overview:** (a) Given a sequence of multiview images, we first run a 3D Face Tracker [Wu et al. 2016] to obtain tracked meshes in consistent topology. (b) Next, we define a patch layout and compute patch centers, together with orientations and positions in the world space (given by a TBNP matrix  $\mathcal{T}_p$ ). This gives us a per-patch coordinate frame, which we combine with individual patch blendweights  $\beta_p$  (obtained from the PBS model Sec. 4.2) for the next steps. (c) Finally, Scaffold-GS anchors  $A_p$  are attached to the patches. Each anchor’s attributes, i.e. position  $\mu$ , scale  $s$ , opacity  $\alpha$ , and anchor feature  $\mathbf{f}$  and optimized together with the global expression MLP  $G$ , per-patch expression MLPs  $\{P_0 \dots P_{p-1}\}$  and scaffold MLPs  $\{\mathcal{F}_\mu, \mathcal{F}_\alpha, \mathcal{F}_q, \mathcal{F}_s, \mathcal{F}_c^{1:P}\}$  for decoding the Gaussian features. Note that the Gaussian primitives (position, scale, and orientation) are predicted in the local coordinate frame of the patch and are deformed to global space with the tracked mesh. The result is a high-quality, re-animatable 3DGS-based avatar.

the following attributes:

$$A_{p,i} = \{\mu_{p,i}, s_{p,i}, \alpha_{p,i}\} \cup \{\mathbf{f}_{p,i}\}, \quad (11)$$

with  $\mu_{p,i}$  denoting anchor’s position in local space of the patch,  $s_{p,i}$  the anchor’s scale, and  $\alpha_{p,i}$  the anchor’s opacity. Similar to Scaffold-GS, we learn per-anchor features  $\mathbf{f}_{p,i}$  containing semantic information describing the behaviour of anchor primitives.

Let  $M_p^{\text{loc}} = \{\mu_{p,1}, \mu_{p,2}, \dots, \mu_{p,N}\}$  denote the set of anchors positions for patch  $p$  such that each anchor is represented in the local coordinate space of the patch. The global positions of the anchors, denoted as  $M_p^{\text{glob}}$  are given by transforming the anchors as follows:

$$M_p^{\text{glob}} = \{\mu_{p,1}^{\text{glob}}, \mu_{p,2}^{\text{glob}}, \dots, \mu_{p,N}^{\text{glob}}\} \text{ with } \mu_{p,i}^{\text{glob}} = \mathcal{T}_p \cdot \begin{bmatrix} \mu_{p,i} \\ 1 \end{bmatrix}. \quad (12)$$

Note that the anchors  $A_{p,i}$  do not define a separate rotation  $\mathbf{R}$ . Empirically, we found that learning separate rotation for anchors did not help with the quality, hence we omit it and use the rotation of the parent patch  $p$  extracted from  $\mathcal{T}_p$  as the anchor rotation.

The number of anchors per patch is dynamically optimized during training. New anchors can be added (densification) or old anchors removed (pruning). Each anchor and their attributes are then used to spawn 3D Gaussian primitives in the global space, which we explain in the next section.

### 5.3 Spawning 3D Gaussian Primitives from Blendweights

One of our core ideas is to leverage per patch expression blendweights  $\{\beta_1, \dots, \beta_p\}$  to synthesizing 3D Gaussian primitives. However, directly using patch blendweights leads to suboptimal results. Therefore, we leverage a set of small per-patch blendweight MLPs  $\mathcal{P} = \{P_1, \dots, P_p\}$  to map patch-specific blendweights to latent expression features. Each patch expression MLP  $P_i$  processes the blendweights  $\beta_i$  of the corresponding patch and predicts latent features  $\mathbf{e}_p$  capturing the localized patch expression information as:

$$\mathbf{e}_p = \mathbf{P}(\beta_p), \quad \forall p \in \{1, 2, \dots, P\}. \quad (13)$$

The per patch blendweight MLPs ensure accurate handling of fine-grained expression details across patches for subsequent processing by Scaffold MLPs. We additionally learn a global expression MLP  $G$  with the concatenated patch blendweights as input for encoding global expression semantics:

$$\mathbf{e}_g = G(\beta_1, \dots, \beta_p). \quad (14)$$

Similar to Scaffold-GS, we use a fixed number of  $L = 5$  Gaussian primitives per anchor and use small Scaffold MLPs to learn and compress the 3DGS attributes. These MLPs directly predict the attributes of all primitives in a single forward pass and are conditioned on anchor attributes, latent expression features and viewing direction. Specifically, we learn a set of Gaussian attribute MLPs  $\mathcal{F} = \{\mathcal{F}_\mu, \mathcal{F}_q, \mathcal{F}_s, \mathcal{F}_\alpha, \mathcal{F}_c^{1:P}\}$  one for each of the Gaussian attributes, except color where we use one MLP per patch. Empirically we found

that this resulted in sharper and more accurate colors. Note that in contrast to vanilla Scaffold-GS, we also use an MLP to predict the Gaussian position  $\mu$ . We use a sigmoid as the last activation layer for  $\mathcal{F}_c^p$  and  $\mathcal{F}_\alpha$ , and a softplus for  $\mathcal{F}_s$ . We omit patch-anchor association index for ease of notation in the following equations.

Given anchor feature  $\mathbf{f}$ , patch expression features  $\mathbf{e}_p$  (Eq. 13), global expression feature  $\mathbf{e}_g$  (Eq. 14), viewing direction  $\mathbf{d}_v$ , we predict the attributes for  $L$  Gaussian primitives per global MLP as:

$$\hat{\boldsymbol{\mu}}_{1:L} = \mathcal{F}_\mu(\mathbf{f}_p; \mathbf{e}_p; \mathbf{e}_g; \boldsymbol{\mu}), \quad (15a)$$

$$\hat{\mathbf{s}}_{1:L} = \mathcal{F}_s(\mathbf{f}_p; \mathbf{e}_p; \mathbf{e}_g; \mathbf{d}_v), \quad (15b)$$

$$\mathbf{q}_{1:L} = \mathcal{F}_q(\mathbf{f}_p; \mathbf{e}_p; \mathbf{e}_g; \mathbf{d}_v), \quad (15c)$$

$$\hat{\alpha}_{1:L} = \mathcal{F}_\alpha(\mathbf{f}_p; \mathbf{e}_p; \mathbf{e}_g; \mathbf{d}_v). \quad (15d)$$

The final Gaussian colors are predicted by the associated patch MLP  $\mathcal{F}_c^p$  to which the anchor belongs as:

$$\hat{\mathbf{c}}_{0:L-1} = \mathcal{F}_c^p(\mathbf{f}_p; \mathbf{e}_p; \mathbf{e}_g; \mathbf{d}_v). \quad (16)$$

The final Gaussian scale  $\mathbf{s}_g$  is the anchor scale  $\mathbf{s}$  combined with Gaussian’s local scale  $\hat{\mathbf{s}}$  as

$$\mathbf{s}_g = \hat{\mathbf{s}} \mathbf{s}. \quad (17)$$

Given an anchor’s position in global space  $\boldsymbol{\mu}^{\text{glob}}$ , scaling  $\mathbf{s}$ , rotation matrix  $\mathbf{R}$  and predicted Gaussian’s local position  $\hat{\boldsymbol{\mu}}$ , the final Gaussian’s position in global space is calculated as:

$$\boldsymbol{\mu}_g = \boldsymbol{\mu}^{\text{glob}} + \mathbf{s} \cdot \mathbf{R} \cdot \hat{\boldsymbol{\mu}}. \quad (18)$$

The Gaussian’s rotation quaternion in the global space  $\mathbf{q}$  is predicted directly from rotation MLP  $\mathcal{F}_q$ . Note that during training we only render Gaussians whose parent’s anchors are visible from the given viewpoint and have an opacity  $\alpha$  above the threshold  $\tau$ :

$$\alpha_g = (\alpha > \tau) \cap \hat{\alpha}. \quad (19)$$

The final Gaussian attributes in the global space are passed into the 3DGS rasterizer  $\mathcal{R}$  and the RGB image is rendered as:

$$I_{\text{RGB}} = \mathcal{R}(\boldsymbol{\mu}_g; \mathbf{s}_g; \mathbf{q}; \alpha_g; \mathbf{c}_g). \quad (20)$$

#### 5.4 Optimization and Regularization

Unlike GHA [Xu et al. 2023a], having a fixed number of Gaussian primitives is insufficient for capturing fine facial details. For instance, representing dynamic wrinkles need more Gaussians with dynamically changing color based on expressions compared to a constant mole on the skin. Therefore, we need an adaptive density control strategy that can add and remove anchors based on view-space gradient of the spawned Gaussians.

We found in our experiments that standard position-based densification used by current methods is slower to optimize and does not lead to sufficient densification to render fine facial features. Instead we use a color-based densification strategy that converges much faster and can more accurately reconstruct fine facial details. We follow the densification strategy similar to Scaffold-GS, except that we use screen space color gradient as our heuristic for densification and pruning anchors. Since we bind anchors to the patch centers, we additionally track one more parameter with the anchor attributes  $\mathcal{A}$ , which is the index of its parent patch. This way we also keep track of the anchors attached to each patch and ensure that every patch always has at least one anchor attached.



Fig. 4. Dataset: One frame selected for each participant from the dataset and corresponding zoom-in for some regions. We selected participants with diverse facial geometry like beard and freckles.

We supervise our renderings with L1 and SSIM loss as:

$$\mathcal{L}_{\text{rgb}} = (1 - \lambda) \mathcal{L}_1 + \lambda \mathcal{L}_{\text{SSIM}}. \quad (21)$$

To further enhance perceptual fidelity, especially when zooming into different facial regions, we additionally apply the LPIPS [Zhang et al. 2018] loss on local image patches as:

$$\mathcal{L}_{\text{patch}} = \frac{1}{J} \sum_{j=1}^J \sum_{k=1}^K \mathcal{L}_{\text{LPIPS}}(I_{\text{RGB}}^j, I_{\text{GT}}^j), \quad (22)$$

where  $I_{\text{RGB}}^j$  and  $I_{\text{GT}}^j$  refer to the  $j^{\text{th}}$  local patch regions from the rendered and ground-truth multiview images. We use  $256 \times 256$  patches and sample 16 local patches uniformly for the facial area.

Our anchor-patch association ensures that anchors be spatially correlated with their parent patches and stay close, e.g. an anchor for the mouth should not be associated with a patch from the hair. Similarly, Gaussians should also stay close to their parent anchor. To address this and improve locally consistent motion, we apply soft regularization on the anchor and Gaussian positions as:

$$\mathcal{L}_{\text{xyz}} = \|\boldsymbol{\mu}\|_2 + \|\hat{\boldsymbol{\mu}}\|_2, \quad (23)$$

where  $\boldsymbol{\mu}$  and  $\hat{\boldsymbol{\mu}}$  refers to anchor and Gaussian position in their respective local space.

We further add a regularization similar to Saito *et al.* [2024] to ensure that the scale of the Gaussian primitives stays in a reasonable range, defined as:

$$\mathcal{L}_{\text{scale}} = \begin{cases} \frac{1}{\max(s, 10^{-7})} & \text{if } s < 0.1, \\ (s - 10.0)^2 & \text{if } s > 10.0, \\ 0 & \text{otherwise.} \end{cases} \quad (24)$$

Our final loss function, therefore can be written as:

$$\mathcal{L} = \mathcal{L}_{\text{rgb}} + \lambda_{\text{patch}} \mathcal{L}_{\text{patch}} + \lambda_{\text{xyz}} \mathcal{L}_{\text{xyz}} + \lambda_{\text{scale}} \mathcal{L}_{\text{scale}}. \quad (25)$$

#### 5.5 Progressive Training

Since we want to focus on high frequency facial details in different facial regions, it is important to train with high resolution data to get the details upon zooming. However, upon directly training with high resolution images (3K), the method trains too slow and fails to converge in a reasonable time. Therefore, we first train with 1K images and progressively increase the resolution to 3K. This

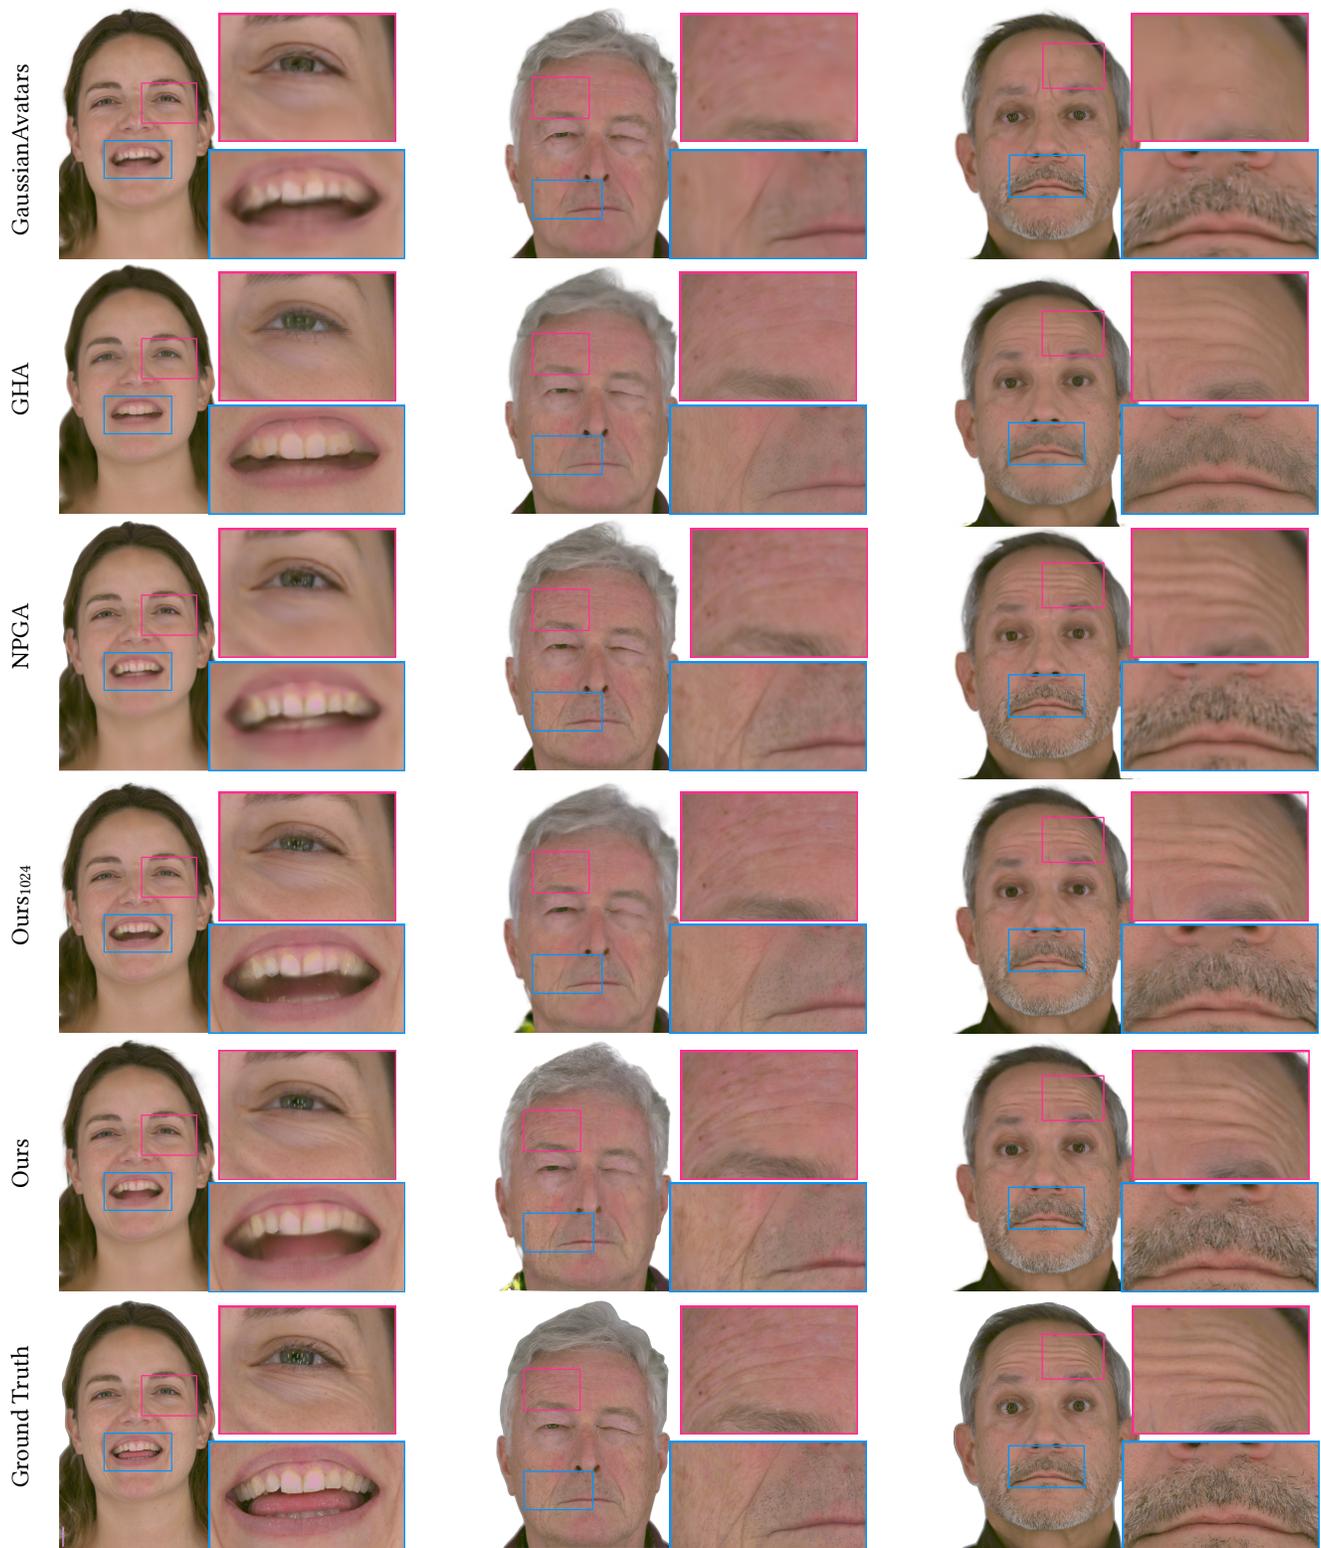


Fig. 5. **Novel View Synthesis:** Our method trained even at 1K resolution outperforms recent state-of-the-art methods and produces significantly sharper renderings. Refining further with 3K resolution data adds even more details. Our method achieves precise reconstruction like wrinkles around the eyes. We request readers to zoom-into the pdf to see subtle differences. Note that all baseline methods were trained at 1K resolution for consistent comparison.

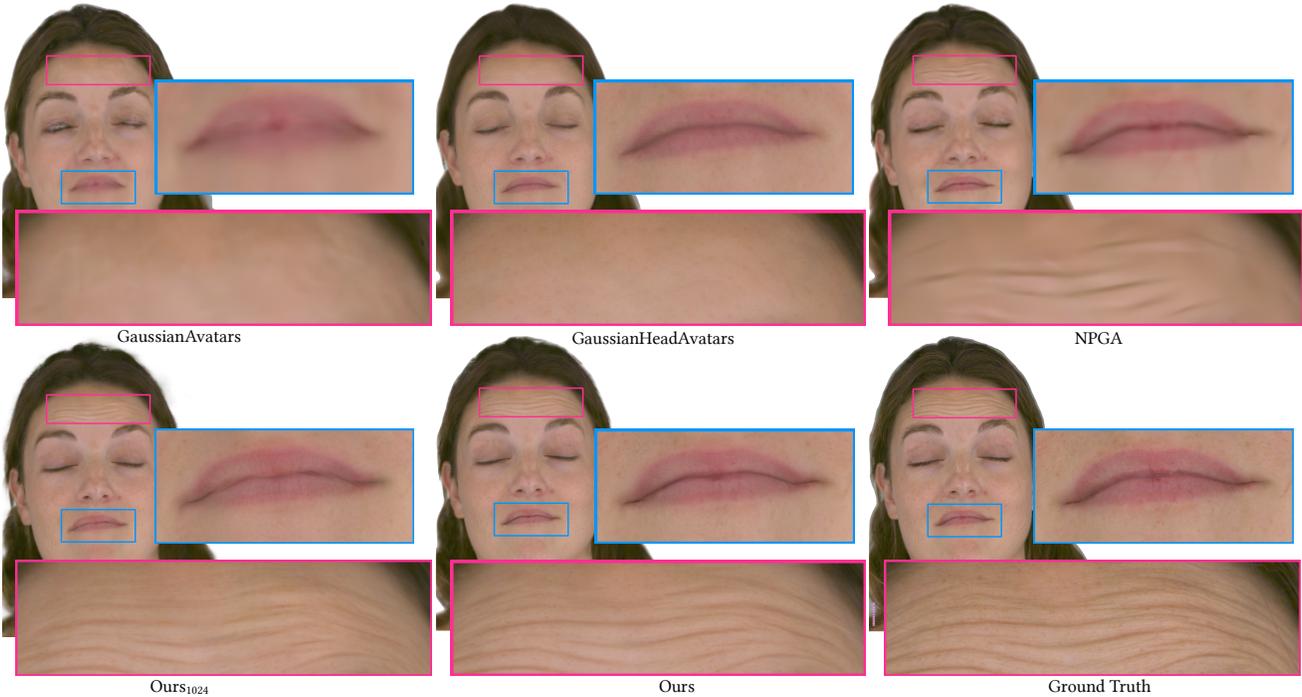


Fig. 6. **Self-Reenactment:** We perform self-reenactment with one held out sequence. Our method achieves ultra-high fidelity results for zoom-ins.

converges much faster while bringing out all the details in different skin areas of the facial region.

## 6 EXPERIMENTS

We collected a high-quality dataset captured using the multi-view setup and pre-processing steps described in Section 4. We selected three challenging participants with facial details like freckles, wrinkles and facial hair. A snapshot of the dataset is shown in Fig. 4. Our model is trained progressively on NVIDIA RTX A6000 (48 GB VRAM) for up to 100,000 iterations (3-4 days) until convergence. On a consumer GPU (NVIDIA RTX 4070Ti 12GB VRAM), it achieves 100.78 FPS for Ours<sub>1024</sub> (our 1K res model) and 76.91 FPS for our full resolution model, making it suitable for interactive applications.

### 6.1 Avatar Reconstruction and Animation

We evaluate ScaffoldAvatar on two tasks, (a) Novel View Synthesis: driving the avatar with training sequences them from held-out/novel viewpoints, and (b) Self-Reenactment: animating the avatar with unseen expressions from held-out viewpoints. For each subject, we downsample images to  $3072 \times 2304$ , and first train at  $1024 \times 768$  and progressively refine at 3K res. We train with all except one workout sequence (used as held out sequence for self-reenactment) and all views except the front view which is used for novel view synthesis.

### 6.2 Metrics

We report the results on a held-out sequence (self-reenactment) and evaluate standard perceptual quality metrics (PSNR, SSIM), and perceptual metric LPIPS [Zhang et al. 2018]. We additionally report results for novel view synthesis (NVS), compare all methods against

Table 1. **Quantitative Comparison.** We compare to recent state-of-the-art baseline methods for novel view synthesis (NVS) and Self-Reenactment.

Method	NVS			Self-Reenactment		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
GA [2024]	26.76	0.9082	0.1489	24.52	0.9078	0.1952
GHA [2023a]	28.93	0.9366	0.1335	26.82	0.9395	0.1911
NPGA [2024b]	30.15	0.9398	0.1312	27.44	0.9335	0.1857
Ours <sub>1024</sub>	32.19	0.9653	0.12597	29.82	0.9515	0.1813
Ours	<b>34.48</b>	<b>0.9712</b>	<b>0.1259</b>	<b>30.37</b>	<b>0.9540</b>	<b>0.1797</b>

held-out camera view from the train set. We use BackgroundMattingV2 [Lin et al. 2020] to remove the background before training. We train all compared baseline methods at 1K resolution and compare against our method trained at both 1K and 3K resolution. Similar to [Giebenhain et al. 2024b], we mask out the torso and neck and report metrics for the facial region.

### 6.3 Results

**6.3.1 Baseline Comparisons.** We compare against recent state-of-the-art methods for NVS and Self-Reenactment. All baseline methods are trained at 1K resolution using the default hyperparameters proposed in their original implementation. For a fair comparison, we also show results of our method trained at 1K. GaussianAvatars [Qian et al. 2024] generates blurry results and cannot synthesize expression dependent details like dynamic wrinkles. GaussianHeadAvatars [Xu et al. 2023a] generates grid-like artifacts due to super-resolution CNN. While NPGA [Giebenhain et al. 2024b] can synthesize dynamic wrinkles and more facial details, it still produces blurry results and artifacts for zoom-ins. Our method outperforms

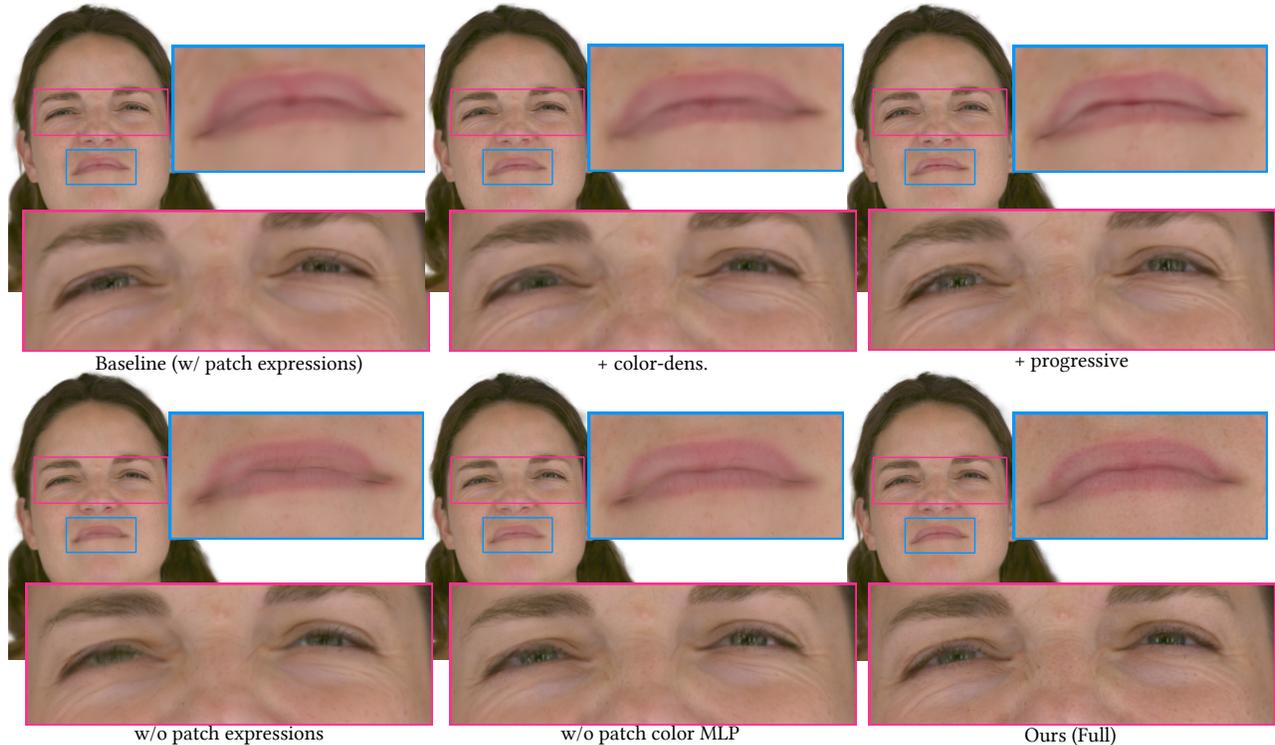


Fig. 7. **Ablation Study:** The naïve baseline of our method "Baseline (w/ patch expressions)" trained with position-based densification and without progressive training produces blurry results. Training with color-based densification (+ color-dens.) improves sharpness and is further refined by progressive training (+progressive). Without patch expressions, our method presents artefacts around lip boundaries and produces relatively blurry results. Without using per-patch color MLP, normal zoom-ins look good but close zoom-ins are blurry. Our full model leveraging all these components achieves the best results.

Table 2. **Ablation Study.** We perform an ablation study on one of the subjects from our dataset and report results for Novel-View Synthesis (NVS) and Self-Reenactment.

Method	NVS			Self-Reenactment		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Baseline (w/ patch expressions)	31.44	0.9635	0.1772	30.34	0.9691	0.1793
+ color-dens	33.76	0.9790	0.1615	32.63	0.9742	0.1656
+ progressive	33.53	0.9779	0.1624	33.24	0.9775	0.1653
w/o patch expressions	34.75	0.9849	0.1291	33.66	0.9822	0.1295
w/o patch color MLP	35.63	0.9837	0.1075	34.39	0.9856	0.1087
<b>Ours (Full)</b>	<b>37.03</b>	<b>0.9880</b>	<b>0.0996</b>	<b>35.145</b>	<b>0.9871</b>	<b>0.1065</b>

these baselines even at 1K resolution. Our full model trained at 3K consistently achieves better results over all the baselines on both tasks perceptually as well as quantitatively. Qualitative results for NVS are shown in Fig. 5 and for Self-Reenactment in Fig. 6. Quantitative results averaged over all subjects are shown in Tab. 1.

**6.3.2 Ablation Study.** We ablate different design choices of our method on one of the subjects from our dataset. We document these results in Tab. 2 and show a visual comparison in Fig. 7. We started out with a "Baseline" version of our method which serves as the naïve baseline of our patch-based avatar. This version is trained with position-based densification, without utilizing progressive training and does not employ per patch color MLP, however it does use per-patch expressions. While this version of our method can produce

light wrinkles, it fails to produce sharp results for fine-scale facial details like freckles on the skin and details on the lips.

**Color-Based Densification.** When changing from view-space positional gradient to view-space color gradient as the heuristic for densifying and pruning the anchors, we notice that our results get sharper. Empirically, we observe that this also converges much faster compared to position based densification.

**Progressive Training.** We use color based densification and progressively increase the resolution of the training images. For the first 20,000 iterations we train with 1K images, then for 35,000 iteration we train with 2K and then finally bump up the resolution to 3K. We notice that progressively increasing the image size also improves the quality of our results compared to training directly at 3K.

**W/o Patch Expressions.** We use color-based densification and progressive training, but we omit using patch-based expressions and instead concatenate all patch blendweights to obtain a single global expression (Eq. 14) which acts as the expression input for our Scaffold MLPs. Note that in this case, we also do not employ per-patch color MLPs but use a single color MLP, since we are operating in global expression space. Although this can produce dynamic wrinkles, we can observe some artefacts around the lips, with intersecting lines between upper and lower lip as shown in Fig. 7.

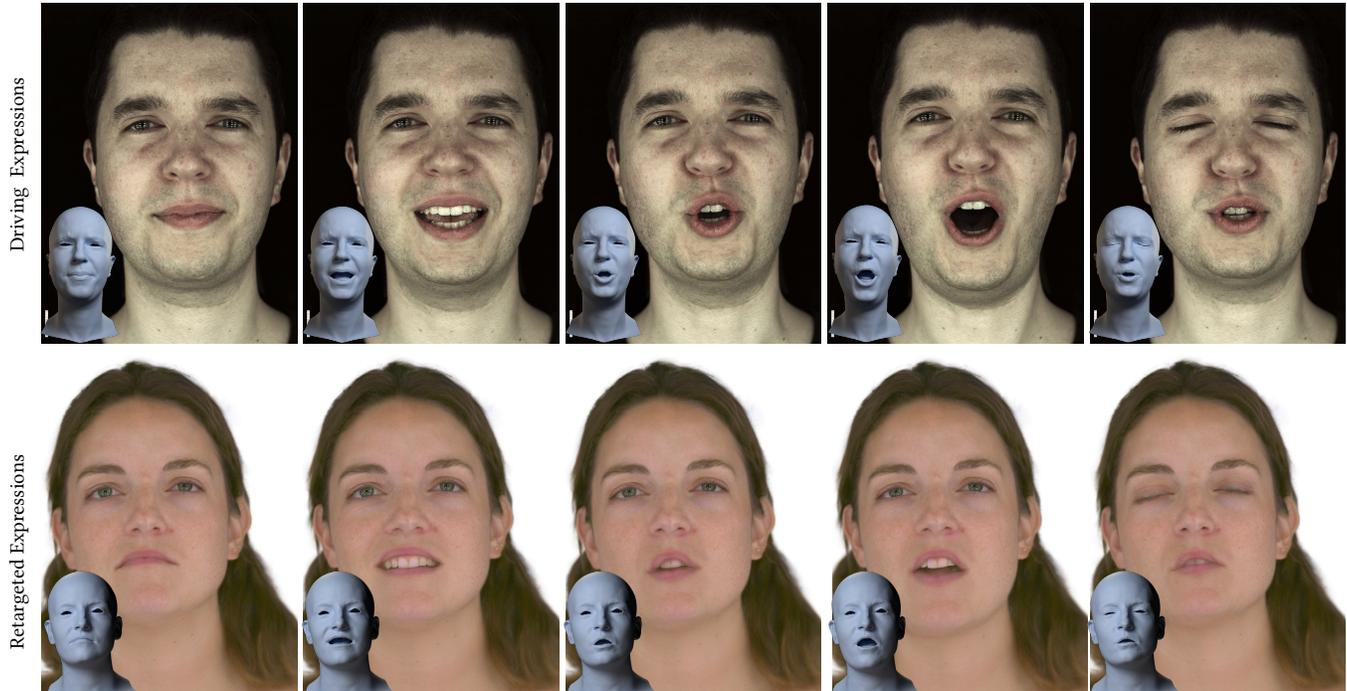


Fig. 8. **Cross-Reenactment:** Qualitative results demonstrating expression transfer from a source actor (top) to our avatar (bottom) across a dialogue sequence. Tracked meshes used in the process are shown in the inset for reference.

*W/o Patch Color MLP.* Finally, we show the importance of using per-patch color MLPs for predicting view and expression dependent color. Although without using per-patch MLPs, our method already achieves good results, the zoom-ins lack sharpness. Per-patch MLPs can handle super-fine facial details much more effectively, as demonstrated in Fig. 7 and reflected quantitatively in Tab. 2.

**6.3.3 Cross-Reenactment.** We reconstruct the 3D facial performance from the driving actor and perform mesh-based retargeting to obtain the mapped geometry performance of our ScaffoldAvatar (using [Chandran et al. 2022]). Then, we fit the patch-blendshape model to this 3D performance to obtain local expression blendweights for driving our avatar with the retargeted geometry. To further improve mouth interior, we added 28 additional patches by leveraging tooth geometry. We train our avatar with expression sequences and evaluate retargeting for a dialogue sequence in Fig. 8.

## 6.4 Limitations

While our method synthesizes high fidelity avatars, it has some limitations. Our template mesh cannot model accessories like eyewear or headwear, and does not include eyes, so eyeball rotation may not be correctly animated. Since we focus on facial microdetails (wrinkles/freckles), we do not track or reconstruct tongue, so certain expressions with tongue motion might not be correctly tracked.

## 7 CONCLUSION

In this work, we propose ScaffoldAvatar, a novel method based on patch blendweights for synthesizing high fidelity 3DGS-based

avatars. The key idea of our method is to leverage per-patch blendweights to learn local deformations on the surface in order to produce more accurate and sharper results. We use anchor-based scaffolding and bind them to the patch centers in order to learn high frequency facial details. We utilize effective densification based on color gradient and patch perceptual loss to obtain sharper results. Our method outperforms current state-of-the-art methods, especially when zoomed into the facial regions. We believe that our method takes an important step towards avatar reconstruction and animation with high fidelity facial details.

## ACKNOWLEDGMENTS

We would like to thank the capture subjects for participating in our dataset. We would also like to thank Loïc Ciccone for help with the Patch Blendweight Optimization codebase, Xinya Ji, Yingyan Xu and Christopher Otto for helpful discussions, Pranjal Singh Rajput for proof-reading and Angela Dai for the video voiceover.

## REFERENCES

- Thabo Beeler, Bernd Bickel, Paul Beardsley, Bob Sumner, and Markus Gross. 2010. High-quality single-shot capture of facial geometry. *ACM Trans. Graph.* 29, 4 (July 2010), 1. <https://doi.org/10.1145/1833351.1778777>
- Volker Blanz and Thomas Vetter. 1999. A morphable model for the synthesis of 3D faces. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99)*. ACM Press/Addison-Wesley Publishing Co., USA, 187–194. <https://doi.org/10.1145/311535.311556>
- Timo Bolkart and Stefanie Wuhrer. 2015. A groupwise multilinear correspondence optimization for 3D faces. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. 3604–3612.
- Alan Brunton, Timo Bolkart, and Stefanie Wuhrer. 2014. Multilinear Wavelets: A Statistical Shape Space for Human Faces. In *Computer Vision – ECCV 2014 (Lecture Notes in Computer Science)*. Springer International Publishing, 297–312.
- Prashanth Chandran, Loic Ciccone, Markus Gross, and Derek Bradley. 2022. Local anatomically-constrained facial performance retargeting. *ACM Trans. Graph.* 41, 4, Article 168 (July 2022), 14 pages. <https://doi.org/10.1145/3528223.3530114>
- Guy Gafni, Justus Thies, Michael Zollhöfer, and Matthias Nießner. 2021. Dynamic Neural Radiance Fields for Monocular 4D Facial Avatar Reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 8649–8658.
- Xuan Gao, Chenglai Zhong, Jun Xiang, Yang Hong, Yudong Guo, and Juyong Zhang. 2022. Reconstructing Personalized Semantic Facial NeRF Models From Monocular Video. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 41, 6 (2022). <https://doi.org/10.1145/3550454.3555501>
- Thomas Gerig, Andreas Morel-Forster, Clemens Blumer, Bernhard Egger, Marcel Lüthi, Sandro Schönborn, and Thomas Vetter. 2017. Morphable Face Models - An Open Framework. arXiv:1709.08398 [cs.CV] <https://arxiv.org/abs/1709.08398>
- Simon Giebenhain, Tobias Kirschstein, Markos Georgopoulos, Martin Rünz, Lourdes Agapito, and Matthias Nießner. 2024a. MonoNPHM: Dynamic Head Reconstruction from Monocular Videos. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Simon Giebenhain, Tobias Kirschstein, Martin Rünz, Lourdes Agapito, and Matthias Nießner. 2024b. NPGA: Neural Parametric Gaussian Avatars. arXiv [cs.CV] (May 2024). arXiv:2405.19331 [cs.CV]
- Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics* 42, 4 (July 2023). <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>
- Tobias Kirschstein, Shenhan Qian, Simon Giebenhain, Tim Walter, and Matthias Nießner. 2023. NeRSemble: Multi-View Radiance Field Reconstruction of Human Heads. *ACM Trans. Graph.* 42, 4, Article 161 (jul 2023), 14 pages. <https://doi.org/10.1145/3592455>
- Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. 2017. Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)* 36, 6 (2017), 194:1–194:17. <https://doi.org/10.1145/3130800.3130813>
- Shanchuan Lin, Andrey Ryabtsev, Soumyadip Sengupta, Brian Curless, Steve Seitz, and Ira Kemelmacher-Shlizerman. 2020. Real-Time High-Resolution Background Matting. arXiv (2020), arXiv–2012.
- Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. 2024. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. openaccess.thecvf.com, 20654–20664.
- Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2021. NeRF: representing scenes as neural radiance fields for view synthesis. *Commun. ACM* 65, 1 (Dec. 2021), 99–106. <https://doi.org/10.1145/3503250>
- Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. 2022. Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM Trans. Graph.* 41, 4, Article 102 (July 2022), 15 pages. <https://doi.org/10.1145/3528223.3530127>
- Thomas Neumann, Kiran Varanasi, Stephan Wenger, Markus Wacker, Marcus Magnor, and Christian Theobalt. 2013. Sparse Localized Deformation Components. *ACM Transactions on Graphics (Proc. of Siggraph Asia)* 32, 6 (Nov. 2013).
- Shenhan Qian, Tobias Kirschstein, Liam Schoneveld, Davide Davoli, Simon Giebenhain, and Matthias Nießner. 2024. Gaussianavatars: Photorealistic head avatars with rigged 3d gaussians. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 20299–20309.
- Shunsuke Saito, Gabriel Schwartz, Tomas Simon, Junxuan Li, and Giljoo Nam. 2024. Relightable Gaussian Codec Avatars. In *CVPR*.
- Kartik Teotia, Hyeonwoo Kim, Pablo Garrido, Marc Habermann, Mohamed Elgharib, and Christian Theobalt. 2024. GaussianHeads: End-to-End Learning of Drivable Gaussian Head Avatars from Coarse-to-fine Representations. *ACM Trans. Graph.* 43, 6, Article 264 (Nov. 2024), 12 pages. <https://doi.org/10.1145/3687927>
- Chenglei Wu, Derek Bradley, Markus Gross, and Thabo Beeler. 2016. An anatomically-constrained local deformation model for monocular face capture. *ACM Trans. Graph.* 35, 4, Article 115 (July 2016), 12 pages. <https://doi.org/10.1145/2897824.2925882>
- Yuelang Xu, Benwang Chen, Zhe Li, Hongwen Zhang, Lizhen Wang, Zerong Zheng, and Yebin Liu. 2023a. Gaussian Head Avatar: Ultra High-fidelity Head Avatar via Dynamic Gaussians. arXiv [cs.CV] (Dec. 2023). arXiv:2312.03029 [cs.CV]
- Yuelang Xu, Lizhen Wang, Xiaochen Zhao, Hongwen Zhang, and Yebin Liu. 2023b. AvatarMAV: Fast 3D Head Avatar Reconstruction Using Motion-Aware Neural Voxels. In *ACM SIGGRAPH 2023 Conference Proceedings*.
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *CVPR*.
- Yufeng Zheng, Wang Yifan, Gordon Wetzstein, Michael J. Black, and Otmar Hilliges. 2023. PointAvatar: Deformable Point-based Head Avatars from Videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wojciech Zielonka, Timo Bolkart, and Justus Thies. 2022. Instant Volumetric Head Avatars. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022), 4574–4584.