

FlashAvatar: High-fidelity Head Avatar with Efficient Gaussian Embedding

Jun Xiang Xuan Gao Yudong Guo Juyong Zhang*
 University of Science and Technology of China
 {junxiang@mail., gx2017@mail., yudong@, juyong@}ustc.edu.cn

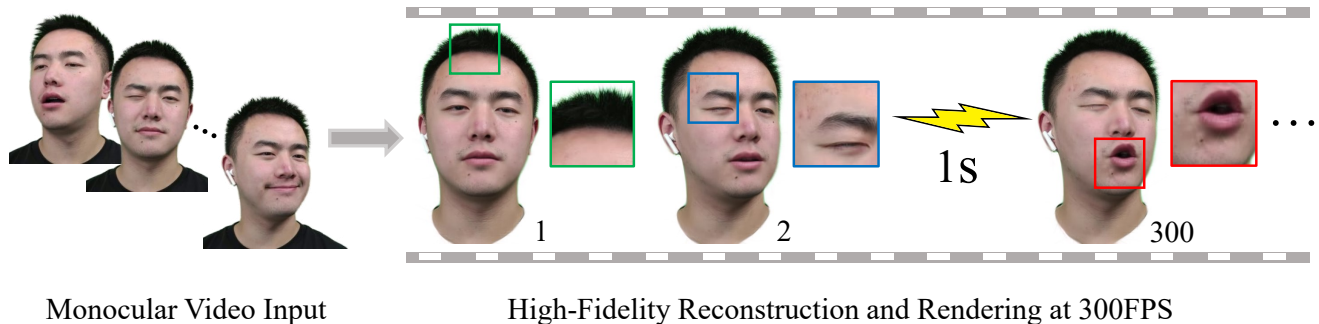


Figure 1. Given a monocular video sequence, our proposed FlashAvatar can reconstruct a high-fidelity digital avatar in minutes which can be animated and rendered over 300FPS at the resolution of 512×512 with an Nvidia RTX 3090.

Abstract

We propose *FlashAvatar*, a novel and lightweight 3D animatable avatar representation that could reconstruct a digital avatar from a short monocular video sequence in minutes and render high-fidelity photo-realistic images at 300FPS on a consumer-grade GPU. To achieve this, we maintain a uniform 3D Gaussian field embedded in the surface of a parametric face model and learn extra spatial offset to model non-surface regions and subtle facial details. While full use of geometric priors can capture high-frequency facial details and preserve exaggerated expressions, proper initialization can help reduce the number of Gaussians, thus enabling super-fast rendering speed. Extensive experimental results demonstrate that *FlashAvatar* outperforms existing works regarding visual quality and personalized details and is almost an order of magnitude faster in rendering speed. Project page: <https://ustc3dv.github.io/FlashAvatar/>

1. Introduction

Achieving low-cost, high-fidelity digital humans with real-time multi-modal interaction, natural expressions and movements, *etc.*, is a key underlying technology for many AR and VR applications, such as immersive remote confer-

encing. With this target in mind, this work aims to present a high-fidelity animatable head avatar that enables efficient reconstruction and lightning-fast rendering, such that the remaining computing resources can support other interactive tasks of multi-modal digital humans.

Previous works have made notable progress, while there still exist some shortcomings. 3D morphable models (3DMMs) [25, 36] based methods [13, 21, 22] are compatible with the standard graphics pipeline and can extrapolate to unseen deformations. However, the limitations of relying on coarse geometry and fixed topology prevent them from modeling complex hairstyles or accessories like eyeglasses. Works [1, 3, 10, 15, 17, 19, 59] building on neural implicit representations [30, 31, 34] could well capture fine features with great rendering quality and 3D consistency but commonly suffer from slow training and inference computation speed. Motivated by works [6, 9, 27, 32, 41] for accelerating Neural Radiance Field (NeRF) [31] rendering, [11, 49, 62] apply voxel representations like voxel grids and multi-level hash tables to speed up head avatar reconstruction. Nevertheless, the volume rendering mechanism of excessive sampling and alpha composition still limits the inference speed.

Recently, 3D Gaussian Splatting (3D-GS) [20] revolutionized radiance field rendering by introducing non-neural 3D Gaussians as geometric primitives and developing a fast rendering algorithm that supports anisotropic splatting. Follow-up works [48, 51] of 3D-GS have already extended it to dynamic scenes by maintaining a canonical Gaussian

*Corresponding Author




Figure 2. Initialization in UV space corresponds to a more uniform Gaussian position distribution, which could model full head details better. We only sample points in the head region, including neck, so the number of sample vertices is smaller than FLAME vertice number 5023.

field and constructing another deformation field conditional on timestamp. However, our experiments have demonstrated that this “canonical + deformation” strategy cannot robustly model dynamic head avatar with complex expressions even if we replace the condition with more meaningful expression code.

Based on these observations, we propose a novel avatar representation named *FlashAvatar*. We initialize a mesh-embedded Gaussian field to model the avatar’s main appearance and facial expressions and learn extra offset to model non-surface features and small facial dynamics. Specifically, we initially attach 3D Gaussians to the mesh surface, which will move along with the mesh. In this way, we do not need to learn large deformations caused by expression changes. However, coarse mesh geometry does not involve non-surface regions like hair or fine facial details like wrinkles. Thus, we use an additional offset network to predict the spatial offsets of 3D Gaussians.

While attaching Gaussians to 3D mesh vertices is a quite straightforward strategy, it is hard to recover complete surface information since the position distribution of vertices is highly uneven. Direct sampling on mesh faces has the same problem of unevenness. Instead, we conduct a flexible UV sampling and turn to maintain a canonical Gaussian field in the UV space. This sampling strategy supports easy density control of Gaussians and generates a much more uniform position distribution (see Fig. 2), which leads to better reconstruction results.

With the help of uniform UV sampling and critical mesh-attached initialization, we achieve photo-realistic head avatar representation with as few 3D Gaussians as possible. Compared with existing 3DMM-based methods, mesh topology will not restrict our representation as tracked meshes only provide initial position distribution and serve as motion-driven tools. Compared to works building on neural implicit representation, we fully introduce geomet-

ric priors, exploit the potential of Gaussian-based radiance field, and thus enable super-fast training and inference. In summary, our contributions include the following aspects:

- We combine Gaussian splats with 3D parametric face model by attaching the Gaussians to the mesh surface and learning extra offsets to model detailed facial dynamics and non-facial features, which leverages dynamic and geometric priors to a great extent and increases the training efficiency.
- Our uniform and flexible UV sampling enables optimal mesh-based initialization, which compresses Gaussian number to 10K level and helps achieve a stable rendering speed at 300FPS at the resolution of 512×512 .
- Experiments demonstrate the high fidelity of our approach even on challenging cases, recovering almost all fine facial details, thin structures, and subtle expressions.

2. Related Work

2.1. Digital Head Model

Digital head model could be classified into explicit and implicit representations. Explicit representations based on mesh have a long history of development. 3DMM [4] first embeds 3D head shape into several low-dimensional PCA spaces. After that, many works [5, 14, 25, 38, 45, 46, 50, 56] are proposed and used for improvement of representation ability. Recently, [22, 43, 44] adopt 2D neural rendering for photo-realistic portrait synthesis but either ignore non-facial regions or suffer from temporal and spatial inconsistencies due to their loose bound to the 3D geometry. [8, 13, 21] opt to learn vertex offset on the head geometry to reconstruct the detailed head model. However, geometry and texture artifacts may occur in hair, eyes, and mouth regions because of the limited representation ability of the mesh model and the approximated differentiable rendering. PointAvatar [60] proposes a deformable point-based representation, which breaks through the limitation of mesh-based models but needs excessive points and long-time training. Implicit head models use neural functions to represent digital head avatars. There have been extensive works on personalized head modeling [1, 10, 58, 59]. They tend to maintain high fidelity but must be more efficient in training or inference. [28] uses volumetric primitives to improve inference efficiency, and [11, 49, 62] use local feature grid to reduce the learning burden of neural network and accelerate the training process. To our knowledge, our work is the first to introduce a mesh-guided Gaussian field for modeling head avatars.

2.2. Scene representations with 3D-GS

3D Gaussian Splatting [20] is currently the SOTA method of scene reconstruction and novel view synthesis regarding rendering speed and visual quality, which inspires a se-

ries of works. [7, 42, 52] adapt 3D-GS into 3D generative tasks by optimizing Gaussian field using score distillation sampling (SDS) [37]. DreamGaussian [42] also designs an efficient mesh extraction algorithm for the Gaussian field. Dynamic3DGS [29] first extends 3D-GS to model dynamic scenes, reconstructing the “point cloud” frame by frame. Different from [29], Deformable3DGS [51] and 4D-GS [48] focus on monocular dynamic scene reconstruction. They both maintain a canonical 3D Gaussian space and optimize an additional deformation field conditional on timestamp. Our work uses 3D-GS to represent dynamic head avatars with complex facial alterations. Rather than adopting the “canonical + deformation” strategy, we attach 3D Gaussians to the head mesh and learn dynamic offsets to model photo-realistic avatars.

2.3. Radiance field acceleration

Neural radiance field (NeRF) [31] and follow-up works [2, 33, 47, 55] significantly develop scene representation but suffer from low rendering efficiency. To accelerate radiance field training and rendering, most works make full use of voxel-based structures like octree [9, 27] and voxel grid [12, 16, 41] by baking information into them which usually needs large cache. INGP [32] adopts a more compressed compact data structure (*i.e.* multi-resolution hash table) and achieves a speedup of several orders of magnitude on training speed but struggles to achieve the visual quality obtained by SOTA NeRF methods [2]. Recently, 3D-GS [20] replaces neural primitives with non-neural 3D Gaussians and designs a fast tile-based rasterizer for Gaussian splats, which guarantees both quality and speed. We apply it to dynamic head representation. Via rational position initialization and density control for Gaussians, we significantly compress the number of used Gaussians and achieve instant training and a stable rendering frame rate at 300FPS.

3. Background

3D Gaussian Splatting. Different from previous methods [24, 53], which use 2D points with normals to represent a scene, 3D-GS [20] chooses 3D Gaussians as geometric primitives of scenes. Every Gaussian is defined by a 3D covariance matrix Σ centered at point μ :

$$g(\mathbf{x}) = e^{-\frac{1}{2}(\mathbf{x}-\mu)^T \Sigma^{-1}(\mathbf{x}-\mu)} \quad (1)$$

To enable differentiable optimization, the positive semi-definite matrix Σ can be decomposed into a rotation matrix \mathbf{R} and a scaling matrix \mathbf{S} corresponding to learnable quaternion \mathbf{r} and scaling vector \mathbf{s} :

$$\Sigma = \mathbf{R}\mathbf{S}\mathbf{S}^T \mathbf{R}^T \quad (2)$$

Given a viewing transformation W and the Jacobian J of the affine approximation of the projective transformation,

3D Gaussians are projected to 2D space for rendering following [63]:

$$\Sigma' = JW\Sigma W^T J^T \quad (3)$$

Besides spatial parameters μ , \mathbf{r} and \mathbf{s} , we attach every 3D Gaussian another two attributes: opacity o and spherical harmonic (SH) coefficients \mathbf{h} representing color \mathbf{c} . The final color for a given pixel is calculated by sorting and blending the overlapped Gaussians:

$$\mathbf{C} = \sum_{i \in N} \mathbf{c}_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j) \quad (4)$$

where α_i represents the density computed by the 2D Gaussian with covariance Σ' multiplied by opacity o .

Analysis. The non-neural nature of 3D-GS reminds us that combining it with concrete mesh will be a new solution to avatar representation. PointAvatar [60] follows similar guidance by using point cloud as the basic representation. In comparison, 3D Gaussian allows anisotropic splatting and fast back-propagation, which is undoubtedly more expressive and easy to optimize.

As in NeRF [31], sampled points near the surface of objects always play a critical role in volume rendering. We assume that modeling avatars with 3D Gaussians follows the same rule, and the ideal Gaussian distribution would be concentrated on the head surface. Thus, it motivates us to attach Gaussians to FLAME mesh surface initially.

The densification scheme of 3D-GS helps model general scenes but leads to explosion and uncertainty of Gaussian’s number, which takes more memory consumption and slows down rendering speed. Since the complexity of head avatars is within a specific range, it is reasonable for us to maintain a fixed number of Gaussians for all subjects instead of adopting the rough splitting strategy of 3D-GS.

4. Methods

Given a monocular video consisting of images $I = \{I_i\}$ along with camera intrinsic parameters \mathbf{K} , camera poses $\mathbf{P} = \{P_i\}$ and tracked FLAME [25] meshes $\mathbf{M} = \{M_i\}$ with corresponding expression codes $\Psi = \{\psi_i\}$, we aim to recover high-fidelity head avatars efficiently with great rendering speed. By fully utilizing the geometric prior knowledge learned in the face-tracking process and the strong representation ability of 3D-GS, we achieve instant training, photo-realistic visual quality, and rendering speed at 300 FPS. An overview of the proposed model is shown in Fig. 3.

4.1. Surface-embedded Gaussian Initialization

Previous head representations based on implicit functions usually build connections with 3DMM by simply utilizing expression code [10] or transformation of the closest point on mesh between canonical and deformed space [1, 62]. In




Figure 3. Overview. We initially maintain the 3D Gaussian field in 2D UV space and embed them into dynamic FLAME mesh surfaces through mesh rasterization. For every surface-embedded 3D Gaussian, the offset network takes tracked expression code and the corresponding position of the Gaussian center on canonical mesh as input, outputs the spatial offset, including position, rotation, and scaling deformation. The deformed Gaussians are then splatted to render the image with a given pose.




Figure 4. To well model interior mouth, we close the mouth cavity of FLAME mesh with additional faces and broaden up corresponding area on UV map.

this way, they fail to fully use the geometric priors of mesh. Our solution is to initially attach 3D Gaussians to the mesh surface, which will move along with the mesh, and we conduct this through UV sampling.

UV Sampling. We conduct UV sampling to locate Gaussian’s position on the mesh surface. By rasterizing the FLAME mesh in world space to UV space, we can get a one-to-one correspondence between UV pixels and mesh surface positions. We sample on the UV map and thus maintain a canonical uniform 3D Gaussian field in 2D UV space. Since the same mesh topology shares fixed UV parameterization, we only need to conduct rasterization [39] once. When expression changes, the corresponding 3D position of Gaussians can be obtained by weighting vertex coordinates using fixed barycentric coordinates.

We can conveniently control Gaussian density by adjusting UV map resolution, sampling interval, and even the covering area of different parts on the UV map based on semantic correspondence. For example, we broaden up the interior mouth area on UV considering the complexity of the internal structure of mouth. It is worth noting that we

add additional faces to close the mouth cavity since original FLAME mesh does not model interior mouth (see Fig. 4).

According to Sec. 3, Gaussian field can be parameterized as $G = \{\mu, \mathbf{r}, \mathbf{s}, o, \mathbf{h}\}$. Through UV sampling, we have defined the initial position of mesh-attached Gaussians μ_M . And in our settings, opacity o , SH coefficients \mathbf{h} , rotation \mathbf{r} and scaling \mathbf{s} are learnable parameters. While the former two attributes, which decide the main appearance of avatars, converge to be fixed, the last two spatial parameters together with μ_M are added with extra deformation to model non-surface features as well as dynamic details of the face.

4.2. Gaussian Offset

We denote the centers of mesh-attached Gaussians as μ_M and corresponding positions on canonical mesh μ_T . Even though main position deformation caused by expression changes has been modeled by μ_M compared to μ_T , non-surface regions and subtle facial details are not considered, and we model them through further adding dynamic spatial offset to Gaussians. The offset network is an MLP F_θ that takes μ_T and ψ as input, and outputs spatial residuals of Gaussians:

$$\Delta\mu_\psi, \Delta\mathbf{r}_\psi, \Delta\mathbf{s}_\psi = F_\theta(\gamma(\mu_T), \psi) \quad (5)$$

where γ denotes the positional encoding as introduced by Mildenhall *et al.* [31]. Then, the final spatial parameters of Gaussians can be computed as:

$$\mu_\psi, \mathbf{r}_\psi, \mathbf{s}_\psi = (\mu_M \oplus \Delta\mu_\psi, \mathbf{r} \oplus \Delta\mathbf{r}_\psi, \mathbf{s} \oplus \Delta\mathbf{s}_\psi) \quad (6)$$

As we initially attach 3D Gaussians to mesh faces, the region a group of Gaussians could influence may expand or shrink with the altering size of mesh faces, especially in the early training process. By adjusting scaling dynamically together with position and rotation, we can better model fixed-size parts like teeth.




Figure 5. Qualitative comparisons with state-of-the-art head avatar reconstruction methods. Our model well reconstructs facial details, thin structures, and subtle expressions while achieving a remarkable rendering speed over 300FPS.

4.3. Training Scheme

Corresponding to expression ψ , our 3D Gaussians field will be $G = \{\mu_\psi, \mathbf{r}_\psi, \mathbf{s}_\psi, o, \mathbf{h}\}$. And following Equation Eq. (4), we will get the rendering image \hat{I} .

To measure the photometric error, we use Huber loss [18] with $\delta = 0.1$:

$$\mathcal{L}_H(x, \hat{x}) = \begin{cases} \frac{1}{2}(x - \hat{x})^2 & \text{if } |x - \hat{x}| < \delta \\ \delta((x - \hat{x}) - \frac{1}{2}\delta) & \text{otherwise} \end{cases} \quad (7)$$

Specifically, we conduct bigger weight for mouth region with mask \mathcal{M} , so the photometric loss \mathcal{L}_C is defined as:

$$\mathcal{L}_C = \mathcal{L}_H(I, \hat{I}) + \lambda_{\text{mouth}} \mathcal{L}_H(I \cdot \mathcal{M}, \hat{I} \cdot \mathcal{M}) \quad (8)$$

In addition to photometric loss \mathcal{L}_C , we adopt perceptual loss $\mathcal{L}_{\text{lipps}}$ proposed in [57] and choose VGG [40] as the backbone of LPIPS. The perceptual loss significantly improves the details of rendered results, and the structure reg-

ularization it brings helps stabilize the training process as well. The total loss is defined as:

$$\mathcal{L} = \mathcal{L}_C + \lambda_{\text{lipps}} \mathcal{L}_{\text{lipps}} \quad (9)$$

4.4. Implementation Details

We implement our network with PyTorch [35], conduct mesh rasterization using PyTorch3D [39] and keep the differential Gaussian rasterization presented by 3D-GS [20]. For FLAME tracking, we use the analysis-by-synthesis-based face tracker from MICA [61] further modified in INSTA [62]. And the expression code ψ is the concatenation of tracked expression coefficients, eyes pose, jaw pose, and eyelids coefficients.

Gaussian initialization and deformation. We set the UV map resolution to 128, sample every UV pixel with correspondence to the head region, including the neck, and the total Gaussian number is 13453. We set the depth of offset MLP $D = 5$ and the dimension of hidden layer $W = 256$.




Figure 6. Our model builds on a non-neural Gaussian field and shows excellent 3D consistency.

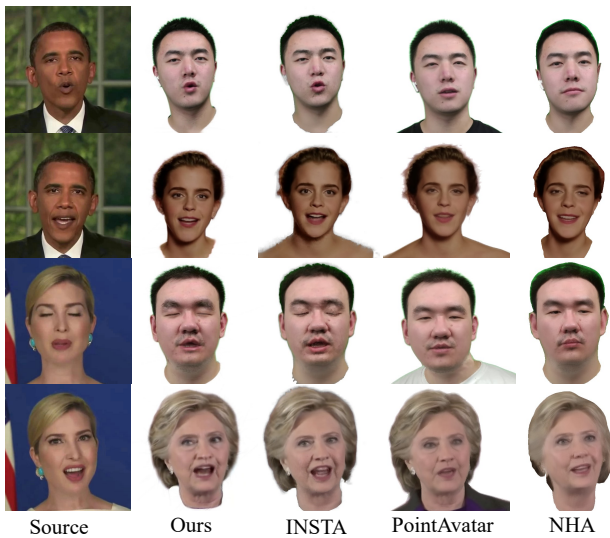


Figure 7. Qualitative results of ours and three other methods on facial reenactment task. Our method preserves personalized facial details in hair, eyes, and interior mouth regions and synthesizes more natural results.

Optimization. Parameters required to be optimized include attributes of 3D Gaussians except for position and parameters of the offset network. We train our models using an Adam optimizer [23] with $\beta = (0.9, 0.999)$. The learning rate of Gaussians’ parameters is the same as the official implementation, while the learning rate of the offset network is $\eta = 1e - 4$. We choose $\lambda_{\text{mouth}} = 40$ and we set λ_{lips} to 0 in the first 15000 training steps and 0.05 later. For each epoch, we randomly sample 2000 frames from the training dataset for training.

5. Experiments

5.1. Dataset

To prove the robustness and fidelity of our methods, We mainly use the data released by previous works [11, 13, 44,

Metrics	NHA	PointAvatar	INSTA	Ours
MSE(10^{-3})↓	1.49	2.47	0.95	0.66
L1(10^{-2})↓	0.99	1.52	0.89	0.83
PSNR↑	28.80	27.03	30.54	32.33
SSIM(10^{-1})↑	9.31	9.00	9.40	9.42
LPIPS(10^{-2})↓	4.01	5.89	3.76	3.23

Table 1. Quantitative comparisons with state-of-the-art head avatar reconstruction methods on public data released by previous works. Our method outperforms others both in pixel-wise error metrics and perceptual quality.

[62], and we appreciate a lot for their sharing. All videos are cropped, sub-sampled to 25 FPS, and resized to 512^2 resolution in advance. The length of the processed video is between 1 and 3 minutes, and we use the last 500 frames as the testing dataset. We use RVM [26] for foreground segmentation and an off-the-shelf face parsing framework [54] for mouth region parsing.

5.2. Comparison with Representative Methods

We compare our method with three representative works, including (1) neural head avatar (NHA) [13], typical work of explicit mesh-based methods; (2) PointAvatar [60], modeling the head geometry with particle-based representation (*i.e.* point clouds) similar to us; and (3) INSTA [62], representative of efficient implicit head representation which creates a surface-embedded dynamic neural radiance field based on neural graphics primitives. Note that for PointAvatar, the full training requires 80GB A100 GPU, but we train it on 32GB V100 and use fewer points and earlier checkpoints exactly following the author’s suggestions. All other experiments were done on 24GB Nvidia RTX 3090. NeRFBlendshape [11], AvatarMAV [49], and INSTA all emphasize training acceleration. We choose INSTA for comparison as it provides tracking code, models neck region, and is almost the latest work among them. FlashAvatar is not only on par with them in training efficiency but also far surpasses them in rendering speed.

Fig. 5 depicts the qualitative comparison between our model and the above methods. As we can see, the representation ability of NHA is restricted by the explicit mesh domain, and it may generate undesired geometric artifacts. INSTA uses neural graphics primitives embedded around the FLAME surface and thus cannot well model accessories like eyeglasses and earphones. Also, it tends to generate smooth results and ignore thin structures, especially in the hair region. As for PointAvatar, the stack of points could recover glasses and earphones, but it still fails to model subtle expressions and clear teeth even with huge memory consumption. In contrast, our method produces photo-realistic images most consistent with the ground truth. We recover




Figure 8. Comparison with “canonical + deformation” strategy. This strategy could get better results with the help of our uniform UV sampling but still fails to capture subtle expression details as well as ours.

almost all fine facial details, thin structures, and subtle expressions with 3D Gaussians in 10K level.

Tab. 1 shows the quantitative comparison between our model and other methods. We compute the average errors of tested videos. The metrics include Mean Squared Error (MSE), L1 distance, PSNR, SSIM, and LPIPS [57].

As both mesh dynamics and later Gaussian deformation condition on tracked expression code disentangled from identity space, we could conduct facial reenactment task at super-fast rendering speed with no difficulty. We show the result of compared methods and ours in Fig. 7. Also, the basic representation of 3D head avatars in our method is pure non-neural 3D Gaussians, so we can freely adjust the global camera pose to generate target results with any desired rendering view (see Fig. 6).

5.3. Comparison with C + D strategy

While the “canonical + deformation” (C + D) strategy is a common way to model dynamics, it struggles to model complex expressions accurately and capture all facial details, especially when we restrict the number of Gaussians to a low level (see Fig. 8). Following former works [48, 51], we randomly initialize the Gaussians in a ball (scaled by the mean size of the head), solely train the canonical 3D Gaussians during the initial 3k iterations and then jointly train Gaussians and the deformation field. However, this common strategy fails to get an acceptable head avatar with many artifacts existing, especially around the head edges. And if we introduce partial geometry priors by initializing canonical Gaussians on the mesh surface the same as ours, most artifacts disappear, but subtle expression details are still not well captured. By comparison, we just need to model extra offset on the basis of a mesh-dependent Gaussian field. Thus, our method can hold exaggerated expressions and preserve fine details with the help of mesh geometry guidance.

5.4. Training Efficiency

We achieve a remarkable rendering speed over 300FPS. Meanwhile, we demonstrate that our training process is super efficient as well in Fig. 9. We are able to recover the coarse appearance of head in several seconds and reconstruct the photo-realistic avatar with fine hair strands and textures within a couple of minutes. We conduct both training and inference on a single Nvidia RTX 3090.

5.5. Ablation Studies

Gaussian Sampling Density. We mainly control the density of Gaussians by adjusting resolutions of the UV map, and Tab. 2 shows the influence of Gaussian sampling density. While sampling more Gaussians will lead to quality improvement, it will also slow down rendering speed. We set UV resolution to 128 but also advise adjusting sampling density according to specific needs.

UV Resolution	PSNR \uparrow	LPIPS(10^{-2}) \downarrow	FPS	# GS
64	30.35	4.47	394	3348
128	30.80	3.47	304	13453
256	31.07	2.99	112	53678

Table 2. Influence of Gaussian density. We set the UV Resolution to 128 in the process of comparison.

Surface Embedding Methods. Attaching Gaussians to mesh vertices cannot converge to satisfactory results (see Fig. 10). Gaussian initialization in UV space is much more uniform than vertice initialization and thus we could get more photo-realistic results with fewer 3D Gaussians.

Distributing Gaussians more carefully or adaptively according to the complexity of different regions and semantic correspondence could get better results, but Tab. 2 has also shown that further processing like local pruning or densification can only get slight improvement on rendering quality and speed on the base of our settings.




Figure 9. Besides significantly fast rendering speed at 300FPS, our training process is also efficient. High-frequency details like hair strands and teeth are fully reconstructed within a few minutes.




Figure 10. More uniform face initialization leads to better results than vertice initialization.




Figure 11. A dynamic offset field is of great importance to modeling fine facial expressions.

Dynamic Offset. Although optimizing a static offset field could well reconstruct static areas like hair regions, it fails to well model facial alterations due to the coarse geometry of FLAME mesh and the complexity of facial expression. As shown in Fig. 11, better visual results with higher fidelity can be obtained by learning a dynamic offset field conditional on expression code.

6. Conclusion and Discussion

In this paper, we have proposed FlashAvatar, which tightly combines a non-neural Gaussian-based radiance field with an explicit parametric face model and takes full advantage of their respective strengths. As a result, it can reconstruct a digital avatar from a monocular video in minutes and ani-

mate it at 300FPS while achieving photo-realistic rendering with full personalized details. Its efficiency, robustness, and representation ability have also been verified by extensive experimental results.

Limitations and Future Work. Our method still has several challenges that need to be addressed in future work. While learning Gaussian offset could compensate for the inaccuracy of tracked mesh surface, our method still relies on a good surface-embedded Gaussian initialization. Therefore, large errors in tracking, especially global pose errors, may cause loss of details or image misalignment. Besides, our representation conditions on tracked expression code and thus cannot model dynamically changing hairs with heavy non-rigid deformation.

Existing works struggle to achieve real-time frame rates for high-fidelity inference, even on high-end hardware. In contrast, FlashAvatar achieves a much faster rendering speed at 300FPS on a consumer-grade GPU with SOTA rendering quality. Therefore, there will be more room for other processes in real-time tasks for multimodal digital humans, such as speech processing, text understanding, and cross-modal translation, with the help of FlashAvatar. One of our future works is to explore its potential in scenarios on mobile and mixed reality devices. We believe that our work is a solid step forward in research and practical applications of multimodal digital humans.

Acknowledgements. This work was supported by the National Natural Science Foundation of China (No. 62122071, No. 62272433) and the Youth Innovation Promotion Association CAS (No. 2018495).

References

- [1] ShahRukh Athar, Zexiang Xu, Kalyan Sunkavalli, Eli Shechtman, and Zhixin Shu. Rignerf: Fully controllable neural 3d portraits. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 20364–20373, 2022. 1, 2, 3
- [2] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. 3
- [3] Alexander Bergman, Petr Kellnhofer, Wang Yifan, Eric Chan, David Lindell, and Gordon Wetzstein. Generative neural articulated radiance fields. *Advances in Neural Information Processing Systems*, 35:19900–19916, 2022. 1
- [4] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pages 187–194, 1999. 2
- [5] Chen Cao, Yanlin Weng, Shun Zhou, Yiying Tong, and Kun Zhou. Facewarehouse: A 3d facial expression database for visual computing. *IEEE Transactions on Visualization and Computer Graphics*, 20(3):413–425, 2013. 2
- [6] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision*, pages 333–350. Springer, 2022. 1
- [7] Zilong Chen, Feng Wang, and Huaping Liu. Text-to-3d using gaussian splatting. *arXiv preprint arXiv:2309.16585*, 2023. 3
- [8] Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. Learning an animatable detailed 3D face model from in-the-wild images. 2021. 2
- [9] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5501–5510, 2022. 1, 3
- [10] Guy Gafni, Justus Thies, Michael Zollhofer, and Matthias Nießner. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8649–8658, 2021. 1, 2, 3
- [11] Xuan Gao, Chenglai Zhong, Jun Xiang, Yang Hong, Yudong Guo, and Juyong Zhang. Reconstructing personalized semantic facial nerf models from monocular video. *ACM Transactions on Graphics (TOG)*, 41(6):1–12, 2022. 1, 2, 6
- [12] Stephan J Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. Fastnerf: High-fidelity neural rendering at 200fps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14346–14355, 2021. 3
- [13] Philip-William Grassal, Malte Prinzler, Titus Leistner, Carsten Rother, Matthias Nießner, and Justus Thies. Neural head avatars from monocular rgb videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18653–18664, 2022. 1, 2, 6
- [14] Yudong Guo, Lin Cai, and Juyong Zhang. 3d face from X: learning face shape from diverse sources. *IEEE Trans. Image Process.*, 30:3815–3827, 2021. 2
- [15] Yudong Guo, Keyu Chen, Sen Liang, Yongjin Liu, Hujun Bao, and Juyong Zhang. Ad-nerf: Audio driven neural radiance fields for talking head synthesis. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 1
- [16] Peter Hedman, Pratul P Srinivasan, Ben Mildenhall, Jonathan T Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5875–5884, 2021. 3
- [17] Yang Hong, Bo Peng, Haiyao Xiao, Ligang Liu, and Juyong Zhang. Headnerf: A real-time nerf-based parametric head model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20374–20384, 2022. 1
- [18] Peter J Huber. Robust estimation of a location parameter. In *Breakthroughs in statistics: Methodology and distribution*, pages 492–518. Springer, 1992. 5
- [19] Boyi Jiang, Yang Hong, Hujun Bao, and Juyong Zhang. Selfrecon: Self reconstruction your digital avatar from monocular video. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1
- [20] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. 1, 2, 3, 5
- [21] Taras Khakhulin, Vanessa Sklyarova, Victor Lempitsky, and Egor Zakharov. Realistic one-shot mesh-based head avatars. In *European Conference on Computer Vision*, pages 345–362. Springer, 2022. 1, 2
- [22] Hyeongwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Niessner, Patrick Pérez, Christian Richardt, Michael Zollhöfer, and Christian Theobalt. Deep video portraits. *ACM transactions on graphics (TOG)*, 37(4):1–14, 2018. 1, 2
- [23] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [24] Georgios Kopanas, Julien Philip, Thomas Leimkühler, and George Drettakis. Point-based neural rendering with per-view optimization. In *Computer Graphics Forum*, pages 29–43. Wiley Online Library, 2021. 3
- [25] Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4d scans. *ACM Trans. Graph.*, 36(6):194–1, 2017. 1, 2, 3
- [26] Shanchuan Lin, Linjie Yang, Imran Saleemi, and Soumyadip Sengupta. Robust high-resolution video matting with temporal guidance. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 238–247, 2022. 6
- [27] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Advances*

- in *Neural Information Processing Systems*, 33:15651–15663, 2020. 1, 3
- [28] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhoefer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering. *ACM Trans. Graph.*, 40(4), 2021. 2
- [29] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In *3DV*, 2024. 3
- [30] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4460–4470, 2019. 1
- [31] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 1, 3, 4
- [32] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. 1, 3
- [33] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11453–11464, 2021. 3
- [34] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019. 1
- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 5
- [36] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3d face model for pose and illumination invariant face recognition. In *2009 sixth IEEE international conference on advanced video and signal based surveillance*, pages 296–301. Ieee, 2009. 1
- [37] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *The Eleventh International Conference on Learning Representations*, 2022. 3
- [38] Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J Black. Generating 3d faces using convolutional mesh autoencoders. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 704–720, 2018. 2
- [39] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv preprint arXiv:2007.08501*, 2020. 4, 5
- [40] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 5
- [41] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5459–5469, 2022. 1, 3
- [42] Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653*, 2023. 3
- [43] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *Acm Transactions on Graphics (TOG)*, 38(4):1–12, 2019. 2
- [44] Justus Thies, Mohamed Elgharib, Ayush Tewari, Christian Theobalt, and Matthias Nießner. Neural voice puppetry: Audio-driven facial reenactment. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI 16*, pages 716–731. Springer, 2020. 2, 6
- [45] Luan Tran and Xiaoming Liu. Nonlinear 3d face morphable model. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7346–7355, 2018. 2
- [46] Daniel Vlasic, Matthew Brand, Hanspeter Pfister, and Jovan Popovic. Face transfer with multilinear models. In *ACM SIGGRAPH 2006 Courses*, pages 24–es. 2006. 2
- [47] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *Advances in Neural Information Processing Systems*, 34:27171–27183, 2021. 3
- [48] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Wang Xinggang. 4d gaussian splatting for real-time dynamic scene rendering. *arXiv preprint arXiv:2310.08528*, 2023. 1, 3, 7
- [49] Yuelang Xu, Lizhen Wang, Xiaochen Zhao, Hongwen Zhang, and Yebin Liu. Avatarmav: Fast 3d head avatar reconstruction using motion-aware neural voxels. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–10, 2023. 1, 2, 6
- [50] Haotian Yang, Hao Zhu, Yanru Wang, Mingkai Huang, Qiu Shen, Ruigang Yang, and Xun Cao. Facescape: a large-scale high quality 3d face dataset and detailed riggable 3d face prediction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 601–610, 2020. 2
- [51] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. *arXiv preprint arXiv:2309.13101*, 2023. 1, 3, 7
- [52] Taoran Yi, Jiemin Fang, Guanjun Wu, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Qi Tian, and Xinggang Wang. Gaussian-dreamer: Fast generation from text to 3d gaussian splatting with point cloud priors. *arxiv:2310.08529*, 2023. 3
- [53] Wang Yifan, Felice Serena, Shihao Wu, Cengiz Öztireli, and Olga Sorkine-Hornung. Differentiable surface splatting

- for point-based geometry processing. *ACM Transactions on Graphics (TOG)*, 38(6):1–14, 2019. [3](#)
- [54] Changqian Yu, Changxin Gao, Jingbo Wang, Gang Yu, Chunhua Shen, and Nong Sang. Bisenet v2: Bilateral network with guided aggregation for real-time semantic segmentation. *International Journal of Computer Vision*, 129:3051–3068, 2021. [6](#)
- [55] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv:2010.07492*, 2020. [3](#)
- [56] Longwen Zhang, Zijun Zhao, Xinzhou Cong, Qixuan Zhang, Shuqi Gu, Yuchong Gao, Rui Zheng, Wei Yang, Lan Xu, and Jingyi Yu. Hack: Learning a parametric head and neck model for high-fidelity animation. *ACM Transactions on Graphics (TOG)*, 42(4):1–20, 2023. [2](#)
- [57] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. [5](#), [7](#)
- [58] Mingwu Zheng, Hongyu Yang, Di Huang, and Liming Chen. Imface: A nonlinear 3d morphable face model with implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20343–20352, 2022. [2](#)
- [59] Yufeng Zheng, Victoria Fernández Abrevaya, Marcel C Bühler, Xu Chen, Michael J Black, and Otmar Hilliges. Im avatar: Implicit morphable head avatars from videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13545–13555, 2022. [1](#), [2](#)
- [60] Yufeng Zheng, Wang Yifan, Gordon Wetzstein, Michael J Black, and Otmar Hilliges. Pointavatar: Deformable point-based head avatars from videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21057–21067, 2023. [2](#), [3](#), [6](#)
- [61] Wojciech Zielonka, Timo Bolkart, and Justus Thies. Towards metrical reconstruction of human faces. In *European Conference on Computer Vision (ECCV)*. Springer International Publishing, 2022. [5](#)
- [62] Wojciech Zielonka, Timo Bolkart, and Justus Thies. Instant volumetric head avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4574–4584, 2023. [1](#), [2](#), [3](#), [5](#), [6](#)
- [63] Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, and Markus Gross. Ewa volume splatting. In *Proceedings Visualization, 2001. VIS'01.*, pages 29–538. IEEE, 2001. [3](#)