Supplementary Material
Coherent 3D Portrait Video Reconstruction via Triplane Fusion

Shengze Wang[1,2], Xueting Li[2], Chao Liu[2], Matthew Chan[2], Michael Stengel[2], Josef Spjut[2], Henry Fuchs[1], Shalini De Mello[2], and Koki Nagano[2]

[1] University of North Carolina at Chapel Hill
[2] NVIDIA

In this supplement, we show additional visual results on in-the-wild (Sec. 1.1) and NeRSemble datasets (Sec. 1.2); provide an explanation of the shoulder pose augmentation process including synthetic multi-view data generation using Next3D [8] (Sec. 2); explain how visibility and occlusion calculations are performed in our method (Sec. 3); visualize the outputs and score matrices that we use to calculate performance metrics (Sec. 4); describe the cropping and training modifications made to the original LP3D (Sec. 5); present three additional sets of quantitative results using different crops of the face (Sec. 6) and, lastly, show how jointly fusing the three planes can cause collapse to 2D (Sec. 7). We strongly recommend that readers **view the accompanying video** with this document for better assessment of the quality of the results of the various methods.

1 Additional Comparisons

In this section, we show more qualitative comparisons between LP3D [9], GPAvatar [3], VIVE3D [4], One-Shot-Avatar [6] and our method. in Figs. 1, 2, 3, 4, and 5. We highly encourage readers to view the **supplementary video**, which provides more visual comparisons.

1.1 In-The-Wild-Data.

In Figs. 1, 2, and 3, we show results of GPAvatar, LP3D and our model on challenging in-the-wild test sequences. Since NeRSemble is a high-quality dataset captured in a controlled studio environment, it is different from real-life usage and limited in terms of lighting conditions, camera viewpoints, and motion blur. Therefore, we capture people of different gender and race in daily environments like offices, apartments, and in outdoor open areas to evaluate the performance of different models in challenging in-the-wild situations. The dataset includes 9 video sequences and 1 image set captured from iPhones, all of which are shown in this supplementary. Our model is able to capture lighting changes (Figs. 1), maintain stable identity (Figs. 3), and remembering the user when their face is partially out of the frame (Figs. 3, second row from the bottom), whereas LP3D shows temporal inconsistency (Figs. 3, red arrows); GPAvatar is not only unable to capture the live lighting condition of the user, but also fails to reconstruct their expressions accurately (Figs. 1 and Figs. 2).

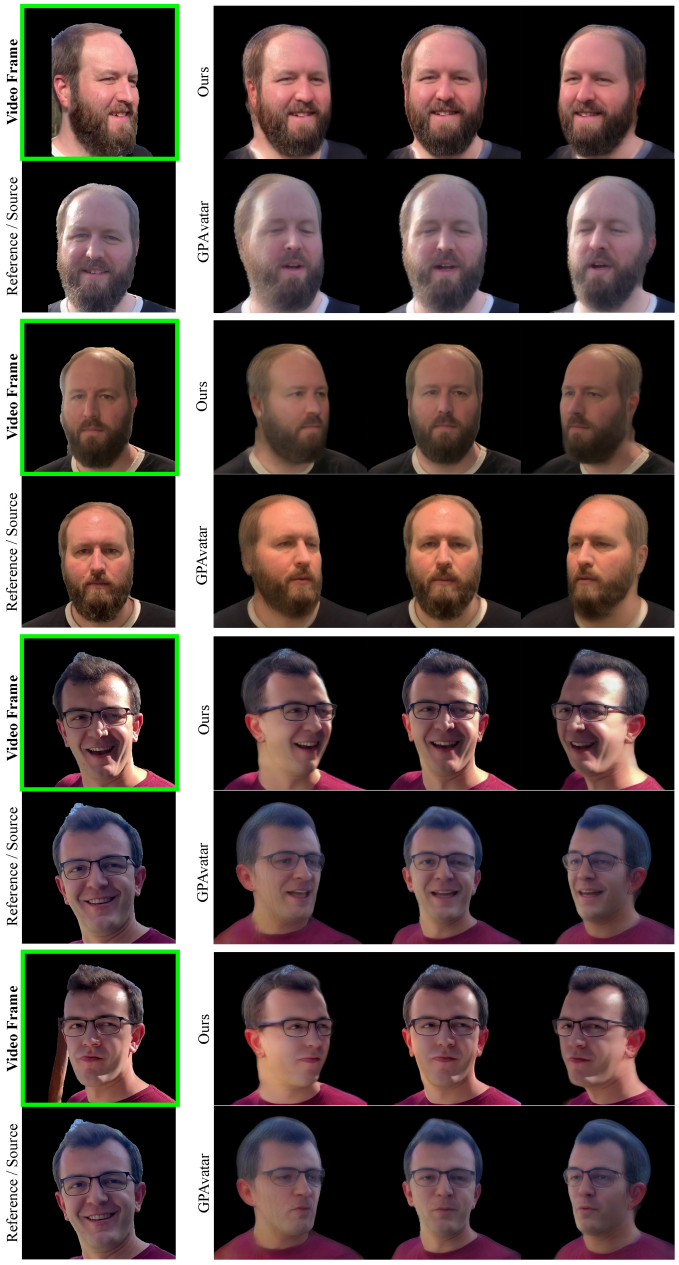

Fig. 1: In-the-wild Lighting (GPAvatar Vs. Ours): Our method captures the dynamic lighting changes in the input video whereas GPAvatar fails to do so. Note that the output of the models should match the lighting and expression of input *Video Frame* (GREEN box).

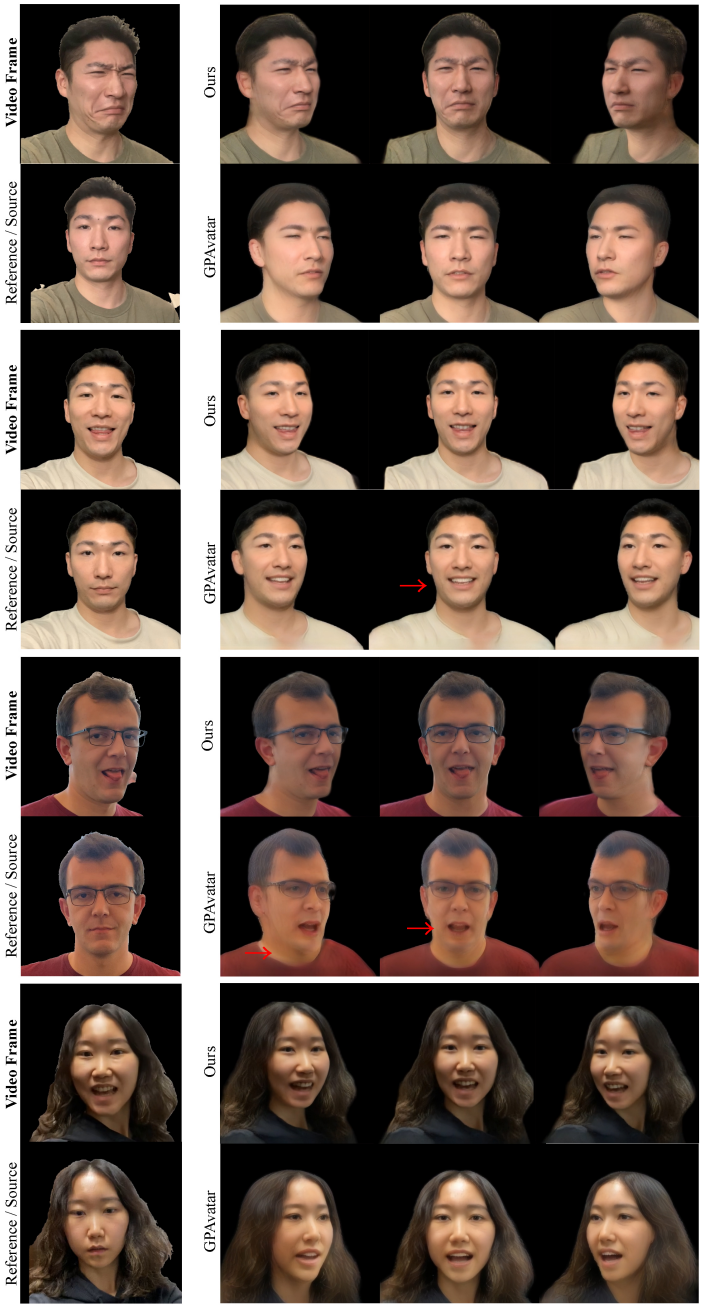

Fig. 2: In-the-wild Expression (GPAvatar Vs. Ours): Our method more accurately captures human expressions in the input video whereas GPAvatar fail to reconstruct authentic expressions. Note that the output of the models should match the lighting and expression of input *Video Frame*.

Fig. 3: In-the-wild Viewpoints (LP3D Vs. Ours): Our method is more robust to variations in the input viewpoint whereas LP3D often performs poorly on rendering novels views that are far from the input view point. Note that the output of the models should match the lighting and expression of input *Video Frame*.

1.2 Additional Results on NeRSemble.

We notice that, despite good numerical performance in terms of LPIPS and PSNR, a closer visual inspection of GPAvatar's results reveals that it is visually not as convincing as the two metrics indicate. This is because it renders dampened expressions (Fig. 4 top examples) and hallucinates parts of the face not present in the reference image (the inner mouth and tongue in Fig. 4 bottom example third row). LP3D is able to reconstruct nuanced facial expressions but struggles to maintain coherent identity when different viewpoints are used as inputs (see Fig. 4 top example first row). Our model achieves both of these properties.

1.3 VIVE3D & Li *et al.* [6]

In our main paper and supplement, we mostly omitted results from Li *et al.* [6] and VIVE3D [4] because of their less competitive results. As mentioned in the main paper, the authors of Li *et al.* [6] kindly performed evaluations for us. Different from other methods, the results are evaluated only on the input viewpoints instead of all 8 viewpoints for NeRSemble. In Fig. 5, we show that this method excels at frontal views but shows significant blurriness from the sides as well as unnatural expressions. On the other hand, VIVE3D is heavily affected by the input viewpoint. It excels at reconstructing the input views but fails to reconstruct other viewpoints well. Comparing to the two methods, we achieve much better reconstruction.

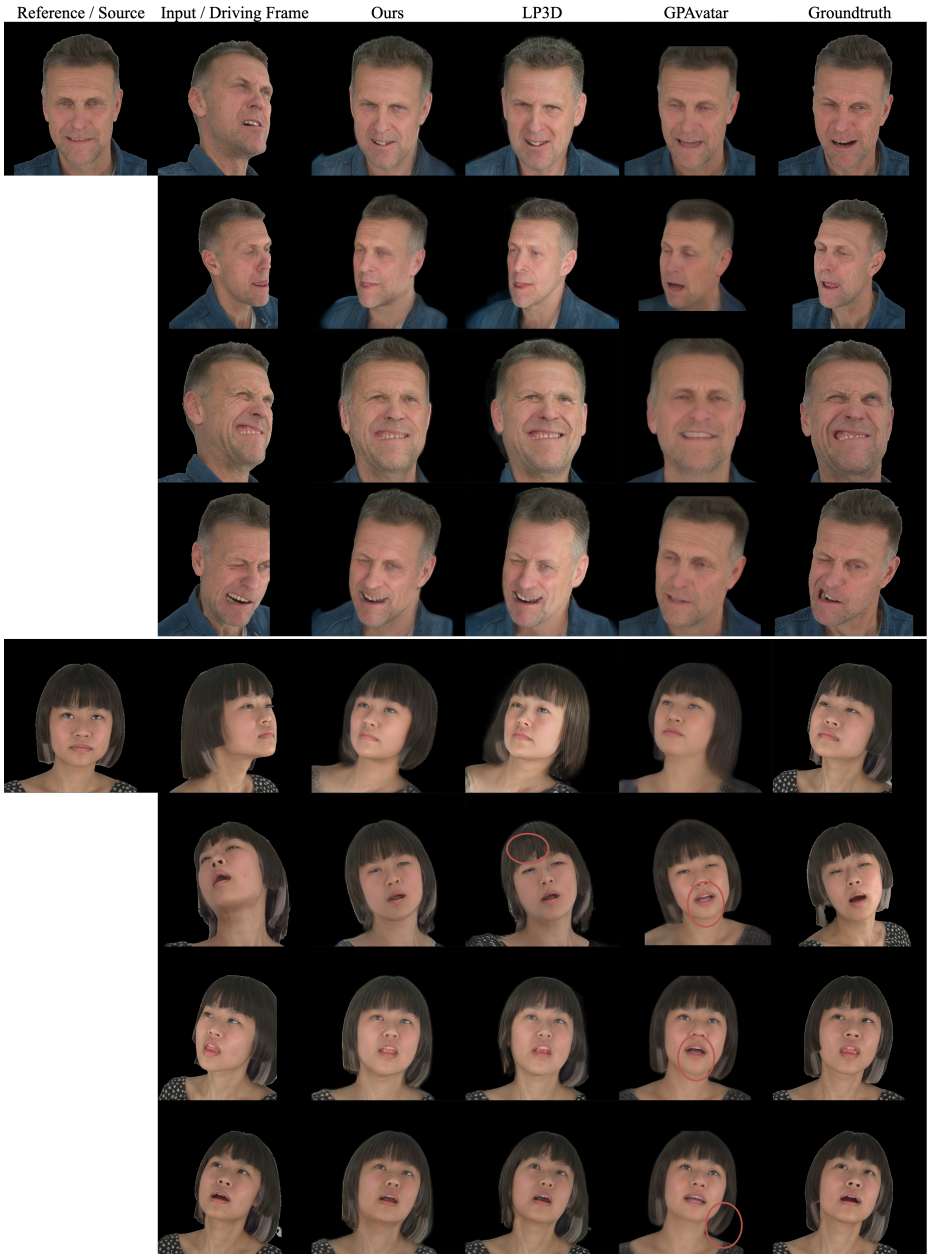

Fig. 4: Example comparisons on NeRSemble sequences. Our model is able to capture extreme expressions and dynamics in hair movement (last row) while maintaining consistent identity despite viewpoint changes. On the other hand, LP3D shows inconsistent identities and GPAvatar exhibits inaccurate expressions and significantly more blurry results. GPAvatar also fails to reconstruct novel content such as the tongue (second last row) and different hair movement (last row). The quality of expression reconstruction is best viewed in the video.

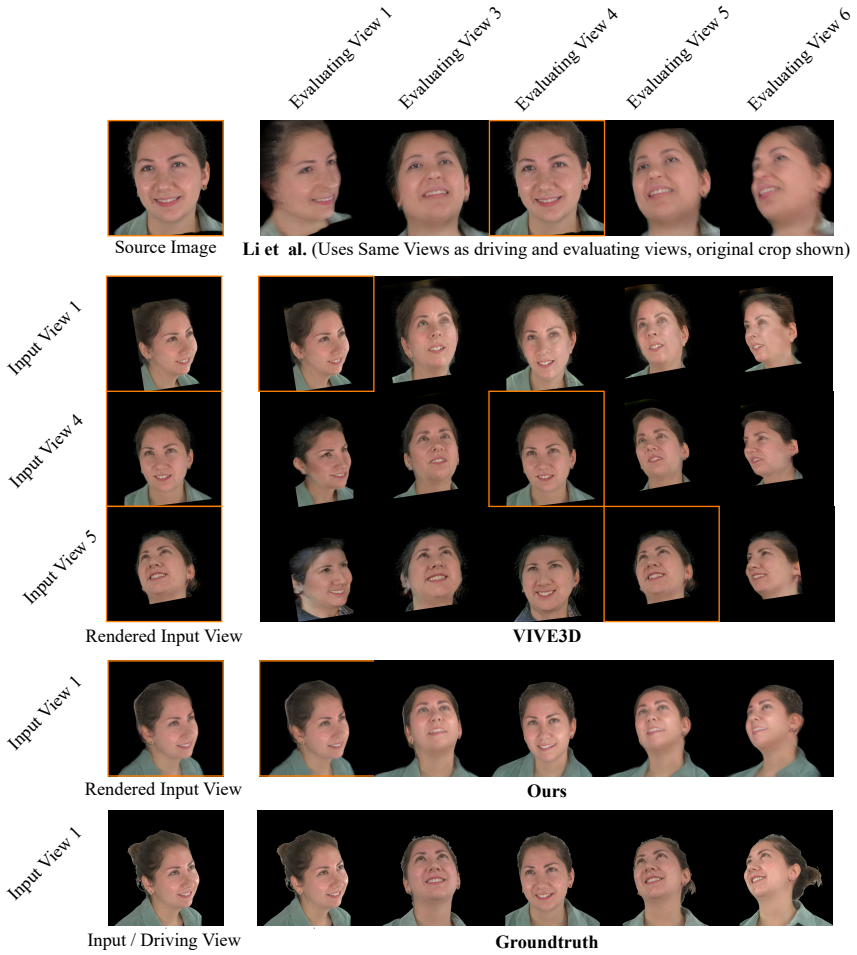

Fig. 5: VIVE3D [4] and Li *et al.* [6]: *"Li* et al. *[6]" Row*: The authors kindly evaluated their methods for us. They evaluated the method on the same driving/input viewpoint (highlighted in orange) instead of all 8 viewpoints. This method excels at frontal views but shows significant blurriness from the sides as well as unnatural expressions. *"VIVE3D [4]" Row*: VIVE3D is heavily affected by the input viewpoint. It excels at reconstructing the input views but fails to reconstruct other viewpoints well. *"Ours" Row*: Our method is able to achieve better reconstructions using the same input view as the other methods. We omit detailed results from VIVE3D and Li *et al.* in the main paper due to their less competitive results. Images shown are at the original resolution.

2 Shoulder Pose Augmentation

As shown in Fig. 6, for training we generate 2 input images (*i.e.*, a Reference Image and an Input Frame in the green box), and 2 groundtruth images using Next3D (in the blue box). We use these images to train our triplane fusion module such that it learns to enhance the reconstruction of the input frame by leveraging a frontal reference frame. When used in practice, the input frame often contains shoulder rotations that are different from that of the reference frame. It is important to reconstruct the varying shoulder pose in the input video because it conveys nuanced body language that is crucial to the perceived realism of an applications such as telepresence.

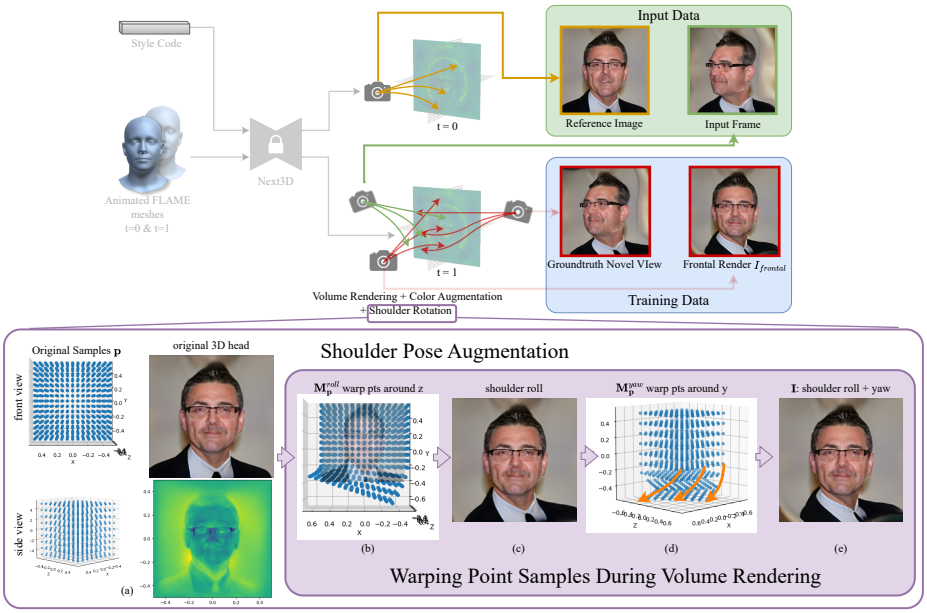

Fig. 6: Shoulder Augmentation. Our data generator (Next3D [8]) does not allow for control over shoulder poses. To enable our model to learn to fuse triplanes with different shoulder poses, we perform shoulder pose augmentation during volume rendering.

We utilize a pretrained 3D GAN, Next3D [8], as our training data generator. However, Next3D does not allow us to synthesize different shoulder poses for the same person. Since it is difficult to change the 3D geometry encoded in triplanes, we synthesize different shoulder poses in the 2D renderings by bending camera rays during volume rendering, *i.e.*, by applying a warping field $\mathbf{M}$ to the 3D points sampled. More formally, we apply the warp fields $\mathbf{M}_{\mathbf{p}}^{roll}$ and $\mathbf{M}_{\mathbf{p}}^{yaw}$ sequentially in order to transform the set of point samples $\mathbf{p}$ used during volume rendering $R(\cdot)$. The final rendered image, I, thus uses the warped point $\mathbf{p}'$ to

sample the triplane T during volume rendering $R(\cdot)$:

$$\mathbf{p}' = \mathbf{M}_{\mathbf{p}}^{yaw}\mathbf{M}_{\mathbf{p}}^{roll}\mathbf{p}, \tag{1}$$

$$\mathbf{I} = R(\mathbf{p}', T). \tag{2}$$

We show an overview of this shoulder augmentation process at the bottom of Fig. 6.

In Fig. 6(a), we show the original 3D head, the Next3D triplane (y-axis upwards, x-axis to the right, z out-of-the-plane), which ranges from -0.5 to 0.5 on all axes, as well as uniform point samples that represent the 3D space before being warped. Then, we warp the camera point samples to achieve shoulder roll (Fig. 6(b)). The warping transform is only applied to the neck and shoulder regions, which are highly consistent in terms of position across Next3D triplanes. This is because 3D GANs like Next3D and EG3D [2] learn a canonical head space from 2D face crops of consistent sizes. Therefore, we find that the neck and shoulder regions can simply be expressed by all point samples $\mathbf{p}_{shoulder} = (x, y, z)$, where $y < y_{chin}$, where $y_{chin} = 0.2$.

We rotate $\mathbf{p}_{shoulder}$ around the top of the neck vertebrae, for which we heuristically use the origin as the rotation pivot. Since a uniform rigid rotation would result in discontinuities, we apply increasingly larger rotations to the points based on their y (vertical) coordinates. Therefore, given a roll rotation angle θ_{base} for the base of the shoulder at $y_{base} = -0.5$, the roll rotation matrix $M_{\mathbf{p}}$ for point $\mathbf{p}$ can be calculated as:

$$d_{chin} = \|y - y_{chin}\|, \tag{3}$$

$$\theta_{\mathbf{p}} = d_{chin}/\|y_{base} - y_{chin}\| \times \theta_{base}, \tag{4}$$

$$\mathbf{M}_{\mathbf{p}}^{roll} = \begin{pmatrix} \cos(\theta_{\mathbf{p}}) & -\sin(\theta_{\mathbf{p}}) & 0 \\ \sin(\theta_{\mathbf{p}}) & \cos(\theta_{\mathbf{p}}) & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{5}$$

Similarly, given the yaw rotation angle ϕ_{base} for the base of the shoulder, the yaw rotation angle $\phi_{\mathbf{p}}$ and matrix $\mathbf{M}_{\mathbf{p}}^{yaw}$ for point $\mathbf{p}$ can be calculated as

$$\phi_{\mathbf{p}} = d_{chin}/\|y_{base} - y_{chin}\| \times \phi_{base}, \tag{6}$$

$$\mathbf{M}_{\mathbf{p}}^{yaw} = \begin{pmatrix} \cos(\phi_{\mathbf{p}}) & 0 & -\sin(\phi_{\mathbf{p}}) \\ 0 & 1 & 0 \\ -\sin(\phi_{\mathbf{p}}) & 0 & \cos(\phi_{\mathbf{p}}) \end{pmatrix}. \tag{7}$$

The final rendered image, I, is thus generated by the volume rendering function $R(\cdot)$ with warped point samples $\mathbf{p}'$ to sample the triplane T using Eqns. (1) and (2).

3 Visibility Estimation and Occlusion Masks

LP3D generates a complete triplane (and thus 3D portrait) from a single image, which inevitably contains occlusion. For example, when the camera captures the

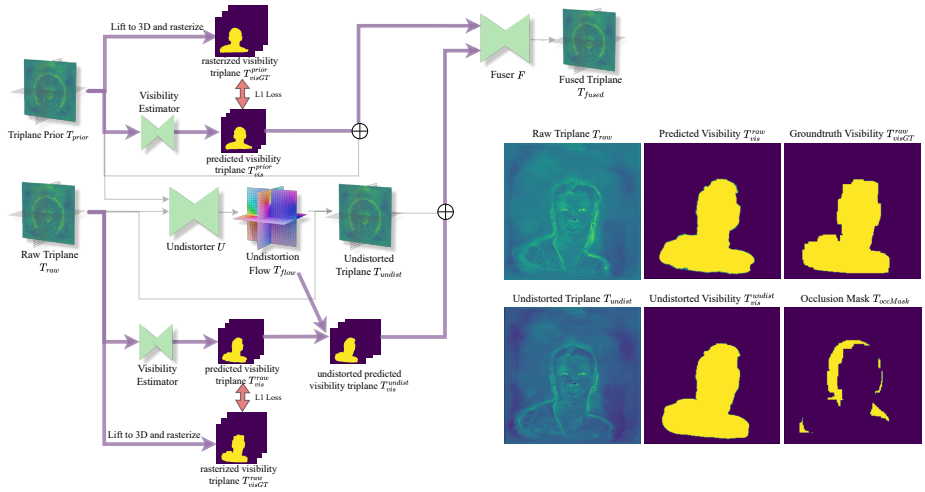

Fig. 7: Visibility Estimation: We show the flow of visibility information in purple.

person from the right, the right side of the face is visible and thus more reliable in the reconstruction whereas the left side of the face is occluded and thus is often inaccurately hallucinated by LP3D. Therefore, to fuse reliable information from the input frame (*i.e.* raw triplane T_{raw}) and the reference image (*i.e.* triplane prior T_{prior}), it is important to inform the fuser F about visible (and thus reliable) regions on the two triplanes.

In Fig. 7, we show how we predict and leverage visibility information by high lighting the data flow of visibility information through our network in purple. First, our model estimates a predicted visibility triplane T^{raw}_{vis} for the raw triplane T_{raw}. Second, the visibility triplane is undistorted alongside T_{undist} using T_{flow}. Finally, the undistorted visibility triplane T^{undist}_{vis} informs the Fuser F about the visibility/reliability of different regions in T_{undist} and allows for better fusion.

Visibility Mask Triplane. There are various ways to compute the visibility information for a triplane. For simplicity, we approximate the actual visibility masks through a rasterization approach: Given a triplane T and its input camera pose C, we generate a pseudo-groundtruth visibility triplane T_{visGT} by first rendering triplane T into a depth map via volume rendering from camera C. We then lift the depth map into a 3D point clouds and rasterize the point cloud back onto the triplane by orthographically projecting the points onto the xy, yx, and xz-planes. The final visibility mask is 1 where points are rasterized and 0 where none is rasterized. Therefore, for triplane prior T_{prior} and the raw triplane T_{raw}, we can calculate pseudo-groundtruth visibility triplanes T^{prior}_{visGT} as well as T^{raw}_{visGT}.

However, this process is expansive due to the volumetric rendering used for depth map generation, we thus develop a Visibility Estimator to directly predict the visibility triplanes. Our Visibility Estimator is a 5-layer ConvNet that predicts visibility maps $T^{prior}_{vis}, T^{raw}_{vis} \in \mathbb{R}^{3\times1\times256\times256}$ from the triplane prior T_{prior} and raw triplane T_{raw}, respectively.

Fig. 8: Occlusion Calculation: The occlusion map is the difference between the approximated foreground mask and the groundtruth visibility triplane T^{raw}_{vis}
.

The two visibility maps are concatenated with T_{prior} and T_{undist} before being input into the Triplane Fuser F. In Fig. 7(right), we show an example of the raw triplane T_{raw} and its predicted visibility triplane T^{raw}_{vis}, undistorted triplane T_{undist}, its visibility triplane T^{undist}_{vis}, and the pseudo-groundtruth visibility triplane T^{undist}_{visGT}.

Occlusion Mask Triplane. In addition to providing the Fuser F with helpful information about visibility, it is also beneficial to emphasize the reconstruction of occluded areas during training because it encourages the model to leverage the frontal reference image for the reconstruction of occluded areas. To achieve this we use an occlusion mask triplane $T_{occMask} \in \mathbb{R}^{3\times1\times256\times256}$ to upweight the triplane loss on occluded areas on the triplane indicated by the mask (see main paper Sec. 4.4). $T_{occMask}$ is calculated as the difference between the visibility triplane T^{raw}_{visGT} of the raw triplane versus the much more complete visibility triplane T^{prior}_{visGT} of the triplane prior.

The Visibility Estimator is supervised via an L_1 loss between the predicted visibility triplane and its groundtruth as:

$$L_{vis} = L_1(T^{raw}_{vis}, T^{raw}_{visGT}) + L_1(T^{prior}_{vis}, T^{prior}_{visGT}). \tag{8}$$

4 Visualization of Score Matrix

In Fig. 10, we show example Score Matrices $\mathbf{S}$ for the NeRSemble dataset's sequence "SEN-10-port_strong_smokey". Each cell $\mathbf{S}_{i,j}$ represents the score of the reconstruction using view i as the input and view j as the evaluation view. Our model achieves higher average as well as more uniform performance, because it has a lower standard deviation and hence more uniform color. Additionally, our model achieves improvements for a majority of the cells (input-evaluation view combinations).

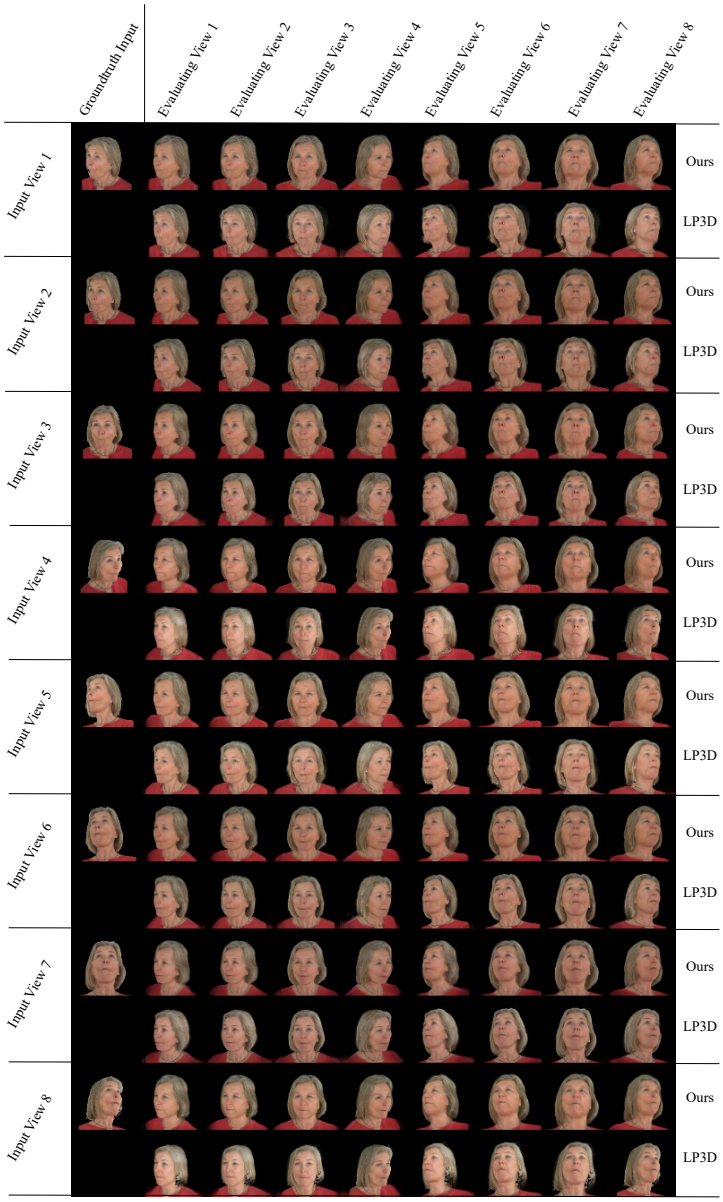

Fig. 9: Output Matrices of LP3D and Our Method: We show an example output matrix from a frame in a NeRSemble test sequence. Each row represents the process of creating a 3D head from the input view (left), and evaluating the reconstruction by rendering on all 8 viewpoints. The images in this $2 \times 8 \times 8$ output matrix are 512×512 each, leading to a large image. The shown output matrix is downsampled for visualization.

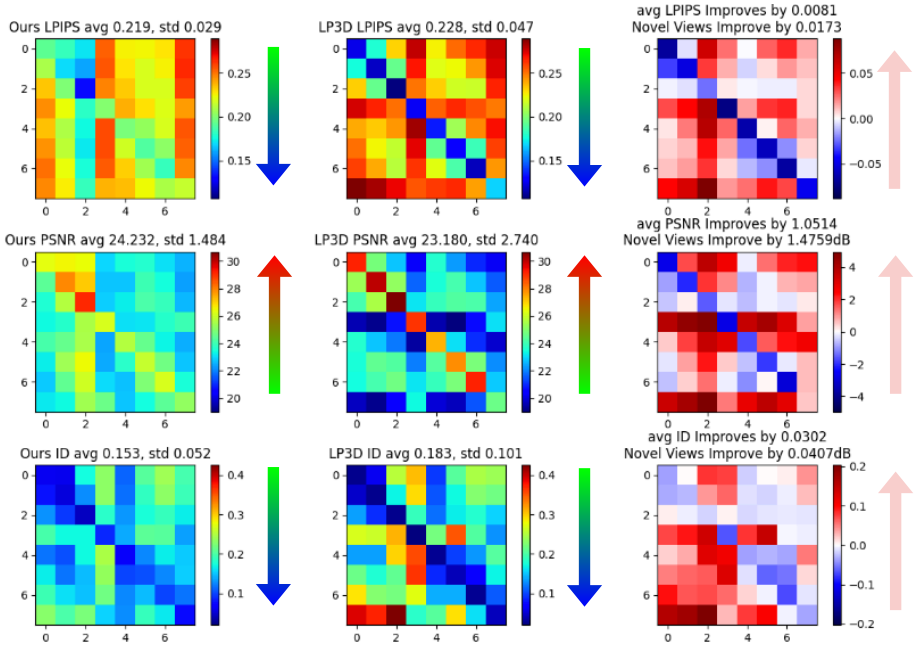

Fig. 10: Score Matrix: We show example Score Matrices **S** for the sequence named "SEN-10-port_strong_smokey". *Left 2 Columns*: Ours' and LP3D's score matrices averaged over the test sequence. LPIPS (top row) and ArcFace ID cosine distance (bottom row) are the lower (greener/bluer) the better, and PSNR (bottom row) is the higher the better (greener/redder)). *Right Column*: red color represents improvement compared to LP3D, and blue represents degradation. Notice that changes in LPIPS and ArcFace ID losses are negated such that positive numbers (red) reflect positive changes. Our model achieves higher average performance as well as more uniform performance (lower standard deviation, more uniform color) whereas LP3D overfits to the input viewpoint and thus achieve higher performance for input views, but performs badly for novel views.

5 Cropping Modifications to LP3D

Our implementation of LP3D mostly follows the original LP3D with a few small modifications. The original LP3D was trained for tight cropping around faces for FFHQ corresponding to the normalized focal length of 4.26 in EG3D [2]. To capture the whole head including shoulders, we increased the field of view and retrained the LP3D with the normalized focal length of 3.12.

6 Performance on Face-only Crops

We use LP3D's face cropping for our model, which includes the face and the shoulders. GPAvatar by default uses center crops (the largest square region at the

center of an image) and do not perform face tracking. This could result in more or less complete reconstructions depending on the image. Due to face cropping inconsistencies among the different methods, their numerical performance can vary based on the kind of cropping used for evaluation. Our model also focuses on shoulders in addition to the head, thus we need to evaluate the models on different input/output image crops for fairness.

"LP3D's Output Crop" rows (Table. 1): For reference, these numbers are also copied from the main paper. Here each of the methods uses its default method to crop the inputs, but we crop the various methods' outputs using LP3D's cropping.

"LP3D's Input Crop" rows: These methods use the same cropped inputs as LP3D instead of applying LP3D's cropping to the output.

"Face Crop" rows: These methods use the same cropped inputs as LP3D, and the rendered images are later cropped around the face region using the face regions detected by the NVIDIA MAXINE AR SDK [1] as on the groundtruth images. This cropping provides the most consistent cropping for all methods but also fails to measure important attributes like shoulder poses and hair.

Since the expression accuracy reported here was calculated using the NVIDIA MAXINE AR SDK [1] on the face crop, the number reported here remains the same as in the main paper and across crops. The ArcFace identity loss is calculated on the entire image and thus varies slightly across different crops.

Our model is the best in expression and identity accuracy among all methods. Despite GPAvatar's good numerical performance on the LPIPS and PSNR metrics, its overall realism is significantly undermined by its dampened expression reconstruction, significant blurriness when viewed from the side, and often inaccurate reconstruction (Fig. 4). Please refer to the supplementary video for more direct visual assessments.

7 Joint vs. Separate Triplane Undistorter and Fuser

As mentioned in the main paper, our Triplane Undistorter and Fuser modules both consist of 3 copies of the same network (with different weights), where each processes one plane in the triplane. It is not unexpected that the Undistorter would require three different instances for the three planes, because it estimates flow maps $T_{flow} \in \mathbb{R}^{3\times2\times256\times256}$ using the optical flow architecture from SPyNet [7]. On the other hand, one might not expect that the Fuser to also require three different instances to process each plane separately. One might expect that jointly fusing the three planes using one transformer allows for communication of information between the 3 planes in a triplane and could thus improve results. However, we find that using a single transformer leads to collapse to 2D (Fig. 11 (left)). We also experimented with first projecting the feature planes into the same feature space before fusion. However, the results remain the same. Whereas using 3 separate smaller networks results in correct fusion (Fig. 11 (right)). We suspect that this is because jointly fusing the triplanes is a significantly more difficult task than fusing each of the planes, separately.

Crop	Method	Expr↓	ID↓	Overall Synthesis Quality		NVS Quality	
				PSNR↑	LPIPS↓	PSNR↑	LPIPS↓
Face	GPAvatar [3]	0.2041	0.2173	21.9434	**0.2327**	**21.9434**	**0.2327**
Crop	LP3D [9]	0.1676	0.1763	21.5092	0.2511	20.7896	0.2670
	Ours	**0.1584**	**0.1644**	**22.1388**	0.2494	21.8849	0.2546
LP3D's	GPAvatar [3]	0.2041	0.2026	22.5624	0.2294	**22.5624**	**0.2294**
Input	LP3D [9]	0.1676	0.2154	22.3309	0.2232	21.5246	0.2374
Crop	Ours	**0.1584**	**0.1865**	**22.7695**	**0.2189**	22.4395	0.2240
LP3D's	GPAvatar [3]	0.2041	0.2074	21.9487	0.2334	21.9487	0.2334
Output	LP3D [9]	0.1676	0.2154	22.3309	0.2232	21.5246	0.2374
Crop	**Ours**	**0.1584**	**0.1865**	**22.7695**	**0.2189**	**22.4395**	**0.2240**

Table 1: Comparison on NeRSemble [5] using face crops: Quantitative performance on the NeRSemble [5] dataset using different input/output cropping settings. The bottom **"LP3D's Output Crop"** rows: These numbers are included in the main paper, where each of the methods use their default methods to crop the inputs. We re-crop their outputs using LP3D's cropping method. The middle **"LP3D's Input Crop"** rows: The methods use the same cropped inputs as LP3D instead of applying LP3D's cropping to their output. The top **"Face"** rows: The methods use the same cropped inputs as LP3D, and the rendered images are cropped around the face region using NVIDIA MAXINE AR SDK's [1] detection. Our method achieves state-of-the-art expression and identity reconstruction across all cropping methods. Please refer to the supplementary video for a better assessment of quality.

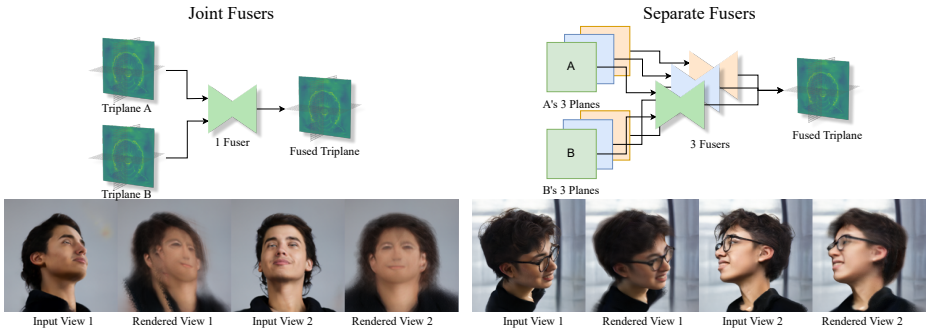

Fig. 11: Joint Fusion Causes Collapse to 2D: *Left*: The joint Fuser treats triplanes as a single feature image of size $96 \times 256 \times 256$. It uses one fusion network to combine two triplanes into a fused one. This approach leads to collapse to 2D as shown at the left-bottom. *Right*: Using Separate Fusers effectively treats each plane in the triplane as a separate entity. Each of the 3 pairs of triplanes are fused separately and combined into the final fused triplane. This approach leads to correct fusion results.

References

1. https://github.com/NVIDIA/MAXINE-AR-SDK
2. Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., Mello, S.D., Gallo, O., Guibas, L., Tremblay, J., Khamis, S., Karras, T., Wetzstein, G.: Efficient geometry-aware 3D generative adversarial networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
3. Chu, X., Li, Y., Zeng, A., Yang, T., Lin, L., Liu, Y., Harada, T.: Gpavatar: Generalizable and precise head avatar from image(s). In: International Conference on Learning Representations (ICLR) (2024)
4. Frühstück, A., Sarafianos, N., Xu, Y., Wonka, P., Tung, T.: VIVE3D: Viewpoint-independent video editing using 3D-Aware GANs. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
5. Kirschstein, T., Qian, S., Giebenhain, S., Walter, T., Nießner, M.: Nersemble: Multi-view radiance field reconstruction of human heads. ACM Trans. Graph. (2023)
6. Li, X., De Mello, S., Liu, S., Nagano, K., Iqbal, U., Kautz, J.: Generalizable one-shot neural head avatar (2023)
7. Ranjan, A., Black, M.J.: Optical flow estimation using a spatial pyramid network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4161–4170 (2017)
8. Sun, J., Wang, X., Wang, L., Li, X., Zhang, Y., Zhang, H., Liu, Y.: Next3d: Generative neural texture rasterization for 3d-aware head avatars. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
9. Trevithick, A., Chan, M., Stengel, M., Chan, E.R., Liu, C., Yu, Z., Khamis, S., Chandraker, M., Ramamoorthi, R., Nagano, K.: Real-time radiance fields for single-image portrait view synthesis. In: ACM Transactions on Graphics (SIGGRAPH) (2023)