

OccFusion: Rendering Occluded Humans with Generative Diffusion Priors

Adam Sun*, Tiange Xiang*, Scott Delp, Li Fei-Fei†, Ehsan Adeli†
Stanford University
{adsun, xtiange}@stanford.edu

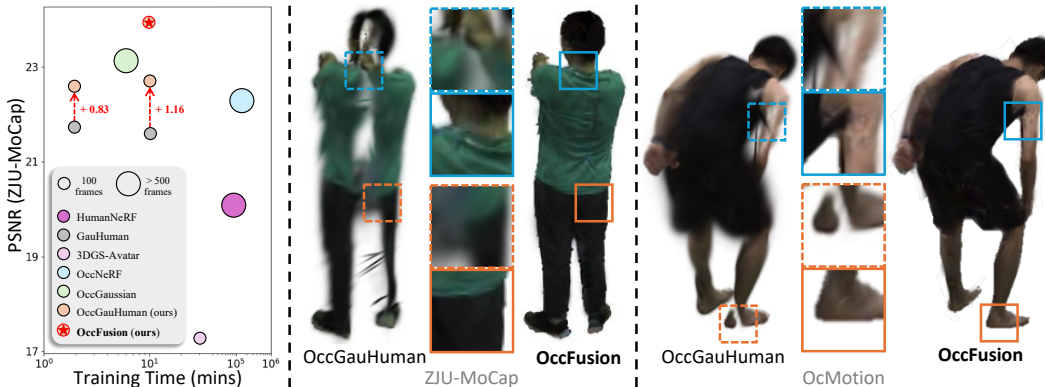


Figure 1: Reconstructing humans from monocular videos frequently fails under occlusion. In this paper, we introduce **OccFusion**, a method that combines 3D Gaussian splatting with 2D diffusion priors for modeling occluded humans. Our method outperforms the state-of-the-art in rendering quality and efficiency, resulting in clean and complete renderings free of artifacts.

Abstract

Most existing human rendering methods require every part of the human to be fully visible throughout the input video. However, this assumption does not hold in real-life settings where obstructions are common, resulting in only partial visibility of the human. Considering this, we present OccFusion, an approach that utilizes efficient 3D Gaussian splatting supervised by pretrained 2D diffusion models for efficient and high-fidelity human rendering. We propose a pipeline consisting of three stages. In the Initialization stage, complete human masks are generated from partial visibility masks. In the Optimization stage, human 3D Gaussians are optimized with additional supervision by Score-Distillation Sampling (SDS) to create a complete geometry of the human. Finally, in the Refinement stage, in-context inpainting is designed to further improve rendering quality on the less observed human body parts. We evaluate OccFusion on ZJU-MoCap and challenging OcMotion sequences and find that it achieves state-of-the-art performance in the rendering of occluded humans. Project page: <https://cs.stanford.edu/~xtiange/projects/occfusion/>.

1 Introduction

Rendering 3D humans from monocular in-the-wild videos has been a persistent challenge, with significant implications in virtual/augmented reality, healthcare, and sports. Given a video of a human

*Equal contribution; junior author listed first.

†Equal mentorship.

moving around a scene, this task involves reconstructing the appearance and geometry of the human, allowing for the rendering of the human from novel views.

When faced with the problem of human reconstruction from monocular video, several works based on neural radiance fields (NeRFs) have achieved promising results [37, 57, 19, 9]. 3D Gaussian splatting [24] further improves upon NeRF-based rendering methods for better performance. By representing the human not as an implicit radiance field but as a set of explicit 3D Gaussians, methods like GauHuman [14] and 3DGS-Avatar [46] are able to render humans comparable in quality to NeRF methods, while taking only a few minutes to train and less than a second to render.

While most human rendering studies assume clean, unobstructed environments, real-world settings like hospitals, stadiums, and construction sites involve frequent occlusions. Current methods struggle in these conditions, often producing artifacts, floating elements, or incomplete body parts. Solutions like OccNeRF [64] and Wild2Avatar [63] attempt to address occlusions but are limited by high computational demands and lengthy training times, making them impractical and restricting their real-world applicability.

In this work, we introduce OccFusion, an efficient yet high quality method for rendering occluded humans. To gain improved training and rendering speed, OccFusion represents the human as a set of 3D Gaussians. Like almost all other human rendering methods [64, 63, 67], OccFusion assumes accurate priors such as human segmentation masks and poses are provided for each frame, which can be obtained with state-of-the-art off-the-shelf estimators such as SAM [25] and HMR 2.0 [8]. However, to ensure complete and high-quality renderings under occlusion, OccFusion proposes to utilize generative diffusion priors, more specifically pose-conditioned Stable Diffusion 1.5 [48] with ControlNet [76] plugins, to aid in the reconstruction process.

Our approach consists of three stages: **(1) *The Initialization Stage***: we utilize segmentation and pose priors to inpaint occluded human visibility masks into complete human occupancy masks to supervise later stages. **(2) *The Optimization Stage***: we initialize a set of 3D Gaussians and optimize them based on observed regions of the human, applying pose-conditioned Score-Distillation Sampling (SDS) to help ensure completeness of the modeled human body in both the posed and canonical space. **(3) *The Refinement Stage***: we utilize pretrained generative models to inpaint unobserved regions of the human with context from partial observations and renderings from the previous stage, further improving the quality of the renderings. Despite taking only 10 minutes to train, our method outperforms the state-of-the-art in rendering humans from occluded videos.

In summary, our contributions are: **(i)** We propose OccFusion, the first method to combine Gaussian splatting with diffusion priors for the rendering of occluded humans from monocular videos. Multiple novel components are proposed along with a three-stage pipeline consisting of Initialization, Optimization, and Refinement stages. **(ii)** We demonstrate that OccFusion achieves state-of-the-art efficiency and rendering quality of occluded humans on both simulated and real-world occlusions.

2 Related Work

2.1 Neural Human Rendering

Traditional methods to reconstruct humans usually require dense arrays of cameras [10, 4, 2] or depth information [71, 50, 4, 5], both of which are unobtainable for in-the-wild scenes. To solve this problem, Neural Radiance Fields (NeRFs) [37] have recently been used to model dynamic humans from monocular videos [57, 9, 19, 17, 72, 52]. These methods achieve high-quality novel view synthesis by parametrizing the human body using an SMPL [35] pose prior and modeling it as a radiance field. However, since NeRFs depend on large Multi-Layer Perceptrons (MLPs), they are computationally expensive, usually taking days to train and minutes to render [24, 14, 46]. To speed up NeRF-based models, multi-resolution hash encoding [40, 43, 6, 18], and generalizability [41, 2, 13, 27] have been proposed. However, these methods either face a rendering bottleneck [14] or an expensive pre-training process, both of which affect their efficiency.

Point-based rendering methods like 3D Gaussian splatting [24] greatly accelerate the rendering of static and dynamic scenes. Recently, there have been an abundance of works applying 3D Gaussian splatting to human rendering tasks [46, 14, 26, 38, 74, 33, 68, 20, 29, 28, 42, 12]. Like NeRF-based approaches, Gaussian splatting-based approaches represent the human in a canonical space

and use Linear Blend Skinning (LBS) to transform the human into the posed space. Gaussian splatting methods achieve state-of-the-art performance of dynamic humans with fast training times and real-time rendering, causing them to be the more desired method [46, 14].

2.2 Occluded Human Rendering

Reconstructing complex scenes in the wild is a well-studied problem. NeRF-W [36] and other works [3, 47, 75] are able to account for photometric variation and transient occluders, allowing them to render consistent representations from unconstrained image collections. However, these works are not designed to handle dynamic objects like humans.

Rendering humans in occluded settings, on the other hand, is relatively understudied. Sun *et al.* [51] utilize a layer-wise scene decoupling model to decouple humans from occluding objects. OccNeRF [64] combines geometric and visibility priors with surface-based rendering to train a human NeRF model, while Wild2Avatar [63] proposes an occlusion-aware scene parametrization scheme to decouple the human from the background and occlusions. While these works provide decent renderings of humans free of occlusions, they are slow and impractical due to their usage of NeRFs. A concurrent work to ours is OccGaussian [67], which also proposes to model occluded humans with 3D Gaussians by performing an occlusion feature query in occluded regions. We provide comparisons to their published results in Table 1.

2.3 Generative Diffusion Priors

Inferring the appearance of unobserved regions of 3D scenes requires the usage of generative models. The recent success of 2D diffusion models has made them the preferred model to use for generation [54, 21, 34, 48, 32]. To lift 2D diffusion models for 3D content generation, DreamFusion [45] proposed Score Distillation Sampling (SDS), a commonly used method for utilizing a pre-trained 2D diffusion model to supervise 3D content generation [30, 70, 53, 55].

Diffusion models can also be used as priors for training NeRFs and Gaussian splatting, combining reconstruction with generation [61, 78, 79, 62, 66, 73]. ReconFusion [61] uses SDS in conjunction with multi-view conditioning to synthesize the appearance of unobserved regions of a scene from sparse views, while BAGS [78] utilizes SDS to supervise a Gaussian splatting model.

3 Preliminaries

Before introducing our method, we provide an overview of key fundamentals in 3D human modeling using SMPL (subsection 3.1). Then, we discuss 3D Gaussian splatting, and how it can be applied to human modeling (subsection 3.2). Finally, we propose OccGauHuman, a simple improvement of GauHuman [14] that is better designed for occluded human rendering (subsection 3.3).

3.1 3D Human Modeling

SMPL [35] is a model that parametrizes the human body with a 3D surface mesh. To transform between the canonical space to a pose space, the Linear Blend Skinning (LBS) algorithm is used. Given a 3D point \mathbf{x}_c in the canonical space and the shape β and pose θ parameters of the human, a point in the posed space can be calculated as:

$$\mathbf{x}_p = \sum_{k=1}^K w_k (G_k(\mathbf{J}, \theta)\mathbf{x}_c + b_k(\mathbf{J}, \theta, \beta)), \quad (1)$$

where J contains K joint locations, G_k and b_k are the transformation matrix and translation vector, and $w_k \in [0, 1]$ are a set of skinning weights. The SMPL representation is commonly used as a geometric prior for human rendering [64, 63, 46, 14, 57, 19, 72].

3.2 Human Rendering with 3D Gaussian Splatting

3D Gaussian splatting. 3D Gaussian splatting [24] models a scene as a set of 3D Gaussians Π . Each Gaussian is defined by its 3D location \mathbf{p}_i , opacity $o_i \in [0, 1]$, center μ_i , covariance matrix Σ_i , and

spherical harmonic coefficients. The i -th Gaussian is defined as $o_i e^{-\frac{1}{2}(\mathbf{p}-\mu_i)^T \Sigma_i^{-1}(\mathbf{p}-\mu_i)}$. During rendering, these 3D Gaussians are mapped from the 3D world space and projected to the 2D image space via α -blending, with the color of each pixel being calculated across the N 3D Gaussians as:

$$C = \sum_{j=1}^N c_j \alpha_j \prod_{k=1}^{j-1} (1 - \alpha_k), \quad (2)$$

where c_j is the color and α is the z -depth ordered opacity. During the training process, 3D Gaussians are adaptively controlled via densification (splitting and cloning) and pruning until they achieve the optimal density to adequately represent the scene.

GauHuman [14]. In the line of work that uses 3D Gaussian splatting for human rendering [46, 26, 29, 12], GauHuman is a representative approach due to its balance between efficiency and rendering quality. After initializing 3D Gaussians on the vertices of the SMPL mesh, GauHuman learns a representation of the human in canonical space and utilizes LBS to transform each individual Gaussian into the posed space. A pose refinement module $MLP_{\Phi_{pose}}$ and an LBS weight field module $MLP_{\Phi_{lbs}}$ are used to learn the LBS transformation, and a merge operation based on KL divergence is used along with splitting, cloning, and pruning to help the 3D Gaussians reach convergence.

We base our method on GauHuman due to its fast training and state-of-the-art representative ability. GauHuman’s code is distributed under the S-Lab license and can be accessed here.

3.3 OccGauHuman: An Improved Baseline for Occlusion Handling

In common human rendering tasks, videos are captured in a clean environment, with every pixel in the image belonging to either the human or the background. By using a semantic segmentation model such as SAM [25] to preprocess a video, we can train the human rendering model only on pixels labeled as "human". However, occlusions in the videos may lead to sparse observations of the human. As a result, fitting NeRF-based human rendering models on only the visible human pixels results in an incomplete geometry with lots of artifacts [64, 63].

Gaussian splatting-based rendering models [24] are especially suitable for human modeling tasks due to their explicit geometry and point-based representation. In this section, we present three straightforward tweaks of GauHuman [14] to make it perform better on videos with occlusions: (1) Firstly, as discussed above, we train the model on visible human pixels only, ensuring that occlusions do not result in learned sparsity on the human model. (2) We adjust the loss weights to put more weight on the mask loss computed between rendered human occupancy maps and the segmentation masks — we found that this helps learn more crisp human boundaries. (3) We disable the densification and pruning of 3D Gaussians during training — this helps maintain a rather complete human geometry based on the SMPL initialization.

The resulting OccGauHuman model serves as an improved baseline for occluded human reconstruction and as a starting point for our method. Benefits brought by our updates compared to the original GauHuman are presented in Table 1, as well as in Figure 7.

4 OccFusion

In our approach, we train a Gaussian splatting-based human rendering model on the visible pixels of a human. However, recovering occluded content for a dynamically moving human is not trivial — humans are usually in challenging poses, and complex occlusions can cause additional issues. It is also essential to preserve a consistent human appearance and geometry across different frames. Considering these challenges, we propose our method OccFusion in multiple separate stages. In the Initialization stage (section 4.1), we inpaint occluded binary human masks for more reliable geometric guidance. In the Optimization stage (section 4.2), we use the inpainted masks to train a human rendering model based on GauHuman [14] while using Score Distillation Sampling (SDS) constraints on both the posed space and canonical space. In the Refinement stage (section 4.3), we fine-tune the trained model from the Optimization Stage with in-context inpainting to further refine the appearance of the human. An overview of our OccFusion is shown in Figure 2.

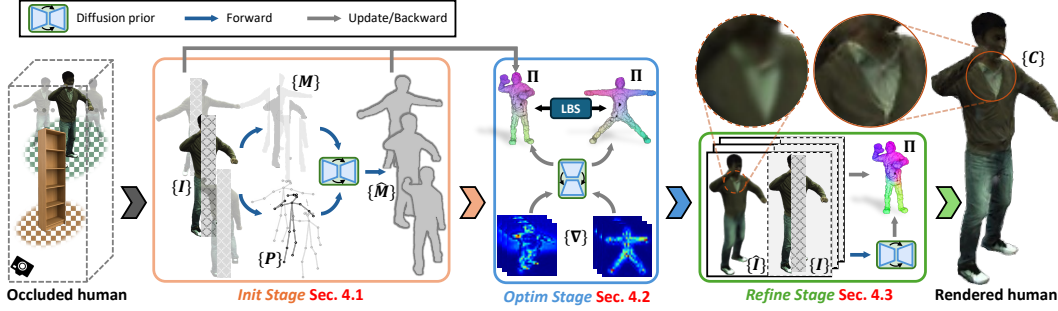


Figure 2: **OccFusion** achieves occluded human rendering via three sequential stages. In the **Initialization Stage**, we recover complete binary human masks $\{\hat{M}\}$ from occluded partial observations $\{I\}$ with the help of segmentation priors $\{M\}$ and pose priors $\{P\}$. $\{\hat{M}\}$ will be further used to help optimize the 3D Gaussians Π in subsequent stages. In the **Optimization Stage**, we apply $\{P\}$ conditioned SDS on both posed human and canonical human to enforce the human occupancy to remain complete. In the **Refinement Stage**, we use the coarse human renderings $\{\hat{I}\}$ from the Optimization Stage to help generate missing RGB values in $\{I\}$ through our proposed in-context inpainting. Through this process, both the appearance and geometry of the human are fine-tuned to be in high fidelity. Training of all three stages takes only **10 minutes** on a single Titan RTX GPU.

4.1 Initialization Stage: Recovering Human Geometry from Partial Observations

Generative diffusion models [48] have demonstrated promise to be used as priors for different tasks [22, 53]. The most straightforward method is to utilize a precomputed segmentation prior M and pose prior P to condition Φ [39, 76] to inpaint $1 - M$ — the image regions that are not occupied by the human. However, there are two significant barriers to such a straightforward approach.




Figure 3: Stable Diffusion 1.5 generations [48] conditioned on a challenging pose P . While conditioning on the original pose results in multiple limbs and other abnormalities, our method of simplifying pose by removing self-occluded joints results in more feasible generations.

Conditioned human generation cannot handle challenging poses. It is true that a conditioned diffusion prior Φ is able to generate detailed images while staying consistent with the condition. However, since diffusion models like Φ are usually overfitted on more commonly seen poses, Φ usually fails to generate reasonable images when conditioned on challenging poses (see Figure 3 middle column). We attribute this limitation to the inappropriate 2D representation of P — when joints occlude each other, it is impossible to tell which joints are closest to the camera when they are projected to 2D. So, we propose to simplify the 2D representation of P . We apply a Z-buffer test on the depth map rendered from the SMPL mesh [35] and then calculate the distance d between its z-axis location and the corresponding 2D z-buffer. Given a pre-defined threshold σ , we deem a joint is self-occluded if $d > \sigma$. Self-occluded joints are ignored when projecting 3D joints onto the 2D canvas for conditioning Φ (see Figure 3 right column). Our simplification improves the generation quality of Φ for challenging poses.

Per-frame inpainting cannot guarantee cross-frame consistency. Compared to image generation models, video generation models [11, 60, 7] are less accessible and much more expensive to run. Without an explicit modeling of object motion in the video, frame-by-frame generation with an image generative model leads to cross-frame inconsistency, which is not desirable for human reconstruction




Figure 4: While generative models provide inconsistent inpainting results, the binary masks that can be extracted from these generated images are much more consistent.

(see Figure 4 middle column). Instead of inpainting the occluded parts of the human directly with Φ , we claim that it is more feasible to inpaint binary human masks since small variations in the human silhouette are more acceptable (see Figure 4 right column). We first inpaint the RGB image \mathbf{I} and then rely on an off-the-shelf segmentation model [23] to obtain the inpainted binary human masks $\{\hat{\mathbf{M}}\}$, which is used to assist the training of the rendering model in subsequent stages.

4.2 Optimization Stage: Enforcing Human Completeness with SDS Regularization

After obtaining the inpainted masks $\{\hat{\mathbf{M}}\}$ that outline a reasonable human silhouette, we build a Gaussian splatting model similar to the one described in section 3.2 for human rendering. The 3D Gaussians Π are initiated as the SMPL mesh vertices, which are able to be deformed to adapt to different poses through SMPL-based LBS (Equation 3.1). With the help of $\{\hat{\mathbf{M}}\}$, the training of Π consists of multiple photometric loss terms \mathcal{L}_{photo} :

$$\lambda_{rgb} L_1(\mathbf{M} \cdot \mathbf{I}, \mathbf{M} \cdot \mathbf{I}') + \lambda_{mask} L_2(\hat{\mathbf{M}}, \mathbf{A}) + \lambda_{sim} \text{SSIM}(\mathbf{M} \cdot \mathbf{I}, \mathbf{M} \cdot \mathbf{I}') + \lambda_{lips} \text{LPIPS}(\mathbf{M} \cdot \mathbf{I}, \mathbf{M} \cdot \mathbf{I}'), \quad (3)$$

where L_1 is the L-1 loss, L_2 is the L-2 loss, $\text{SSIM}(\cdot)$ is the SSIM function [56], LPIPS is the VGG-based perceptual loss [77], \mathbf{I}' is the rendered image from Π , and \mathbf{A} is the rendered human occupancy map. Each of the loss terms is scaled by a weight hyperparameter λ .

Even with the supervision of $\{\hat{\mathbf{M}}\}$, geometry inconsistency still exists. Although inconsistent human masks affect the training of Π much less than inconsistent images, human completeness cannot be guaranteed without further steps.

Using diffusion priors to enforce human completeness. We build off of the insights from [53, 59, 70] and apply Score Distillation Sampling (SDS) [45] to improve the quality of human renderings and reduce artifacts. Instead of applying SDS on RGB images \mathbf{I}' , which causes appearance inconsistency, we apply it directly to the rendered human occupancy maps \mathbf{A} so that diffusion scores are propagated to encourage complete \mathbf{A} :

$$\mathcal{L}_{\text{SDS}}^{(\mathbf{P})} = \mathbb{E}_{t, \epsilon} \left[w(t) (\epsilon_\phi(\mathbf{A}; t, \mathbf{P}) - \epsilon) \frac{\partial \mathbf{A}}{\partial \Pi} \right], \quad (4)$$

where t is a scheduled time stamp, $w(\cdot)$ is a weighting function, $\epsilon(\cdot)$ is the UNet noise estimator in Φ , and ϵ is the injected Gaussian noise.

Using diffusion priors to regularize canonical pose. In-the-wild videos often involve very sparse observations of the human, with only incomplete regions of the human visible in each frame. To further enforce completeness, we propose to render the human in the canonical Da-pose $\hat{\mathbf{P}}$ with the human oriented at a random angle $\in \{k \frac{\pi}{9}, k \in \mathbb{Z}\}$. Applying SDS on the canonical renderings serves as regularization and is randomly activated during training. Overall, at each training step in the Optimization Stage, the 3D Gaussians Π are optimized towards:

$$\nabla_{\Pi} \left[\mathcal{L}_{photo} + \rho \cdot \lambda_{pose} \mathcal{L}_{\text{SDS}}^{(\mathbf{P})} + (1 - \rho) \cdot \lambda_{can} \mathcal{L}_{\text{SDS}}^{(\hat{\mathbf{P}})} \right], \quad (5)$$

where ρ is a random variable that has a 75% chance to be 1 and 0 otherwise. The Optimization stage results in a complete and coherent geometry regardless of the viewing angle.

Table 1: Quantitative comparison on the ZJU-MoCap and OcMotion datasets. LPIPS values are scaled by $\times 1000$. We color cells that have the **best** and **second best** metric values.

Methods	ZJU-MoCap [44]			OcMotion [15]		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR* \uparrow	SSIM* \uparrow	LPIPS* \downarrow
HumanNeRF [57]	20.67 \ddagger	0.9509 \ddagger	-	-	-	-
3DGS-Avatar [46]	17.29 \dagger	0.9410 \dagger	63.25 \dagger	9.788 \dagger	0.7203 \dagger	188.1 \dagger
GauHuman [14]	21.55	0.9430	55.88	15.09	0.8525	107.1
OccNeRF [64]	22.40 \ddagger	0.9562 \ddagger	43.01 \ddagger	15.71	0.8523	82.90
OccGaussian [67]	23.29 \ddagger	0.9482 \ddagger	41.93 \ddagger	-	-	-
Wild2Avatar [63]	-	-	-	14.09 \S	0.8484 \S	93.31 \S
OccGauHuman	22.71	0.9492	54.60	18.85	0.8863	86.53
OccFusion	23.96	0.9548	32.34	18.28	0.8875	82.42

* Metrics calculated on **visible pixels** only.

\dagger Model trained for 5k iterations with $\times 3$ **training time**.

\ddagger Results taken from OccGaussian [67], using $\times 5$ **training frames**.

\S Model trained under the default setting [63] using $\times 2$ **training frames**.

4.3 Refinement Stage: Refining Human Appearance via In-context Inpainting

As shown in Figure 6 Exp. C and D, applying diffusion priors on rendered human occupancy maps is not able to recover the missing appearances of the human. This motivates the need for a subsequent stage that keeps refining Π for better appearance.

The refinement of the appearance of 3D objects is not a new topic [53, 31, 70]. However, no existing generative models are capable of handling the consistency of appearance of a human across different frames and poses. We attribute this difficulty to the denoising process used in generative priors — random noise is injected to rendering at each SDS step which leads to uncertain results. This is infeasible for reconstruction tasks, which require frame-consistent representations that agree with all observations.

Our approach focuses on generating inpainted images of the occluded human offline to use as references. We first identify the occluded regions to be inpainted \mathbf{R} by using the rendered human occupancy masks \mathbf{A} from the Optimization Stage and pre-computed human visibility masks \mathbf{M} : $\mathbf{R} = (1 - \mathbf{M}) \cdot \mathbf{A}$. In order to encourage the generated regions to be more consistent with the partial observations, we propose in-context references inspired by in-context learning in language models [1]. Although renderings from the Optimization Stage lack sharp and high-fidelity details, they resemble complete human geometries and possess good enough features that can be used as a coarse reference to guide Φ to inpaint similar contents at occluded body regions. To achieve this, we stack $\hat{\mathbf{I}}$ and \mathbf{I} together as a single image input to Φ with an additional prompt phrase — "the same person standing in two different rooms".

We use the inpainted RGB images $\{\tilde{\mathbf{I}}\}$ along with other priors to finetune Π via photometric losses. Since diffusion models still tend to be somewhat inconsistent, we smooth training by putting more weight on perceptual loss terms and use L1 loss for the pixel-wise loss terms for its high robustness to variance:

$$\nabla_{\Pi} \left[\lambda_{rgb} L_1(\mathbf{M} \cdot \mathbf{I}, \mathbf{M} \cdot \mathbf{I}') + \lambda_{mask} L_2(\hat{\mathbf{M}}, \mathbf{A}) + \lambda_{gen} L_1(\tilde{\mathbf{I}}, \mathbf{R} \cdot \mathbf{I}') + \lambda_{lips} \text{LPIPS}(\mathbf{I}, \mathbf{I}') \right]. \quad (6)$$

We train our entire pipeline for only **10 minutes** on a single TITAN RTX GPU. More implementation details are provided in supplementary materials.

5 Experiments

In this section, we conduct quantitative and qualitative evaluation of our approach against state-of-the-art methods. Then, we conduct ablation studies of our entire pipeline, demonstrating that each stage is necessary for optimal performance. **More experiments and results can be found in supplementary materials.**

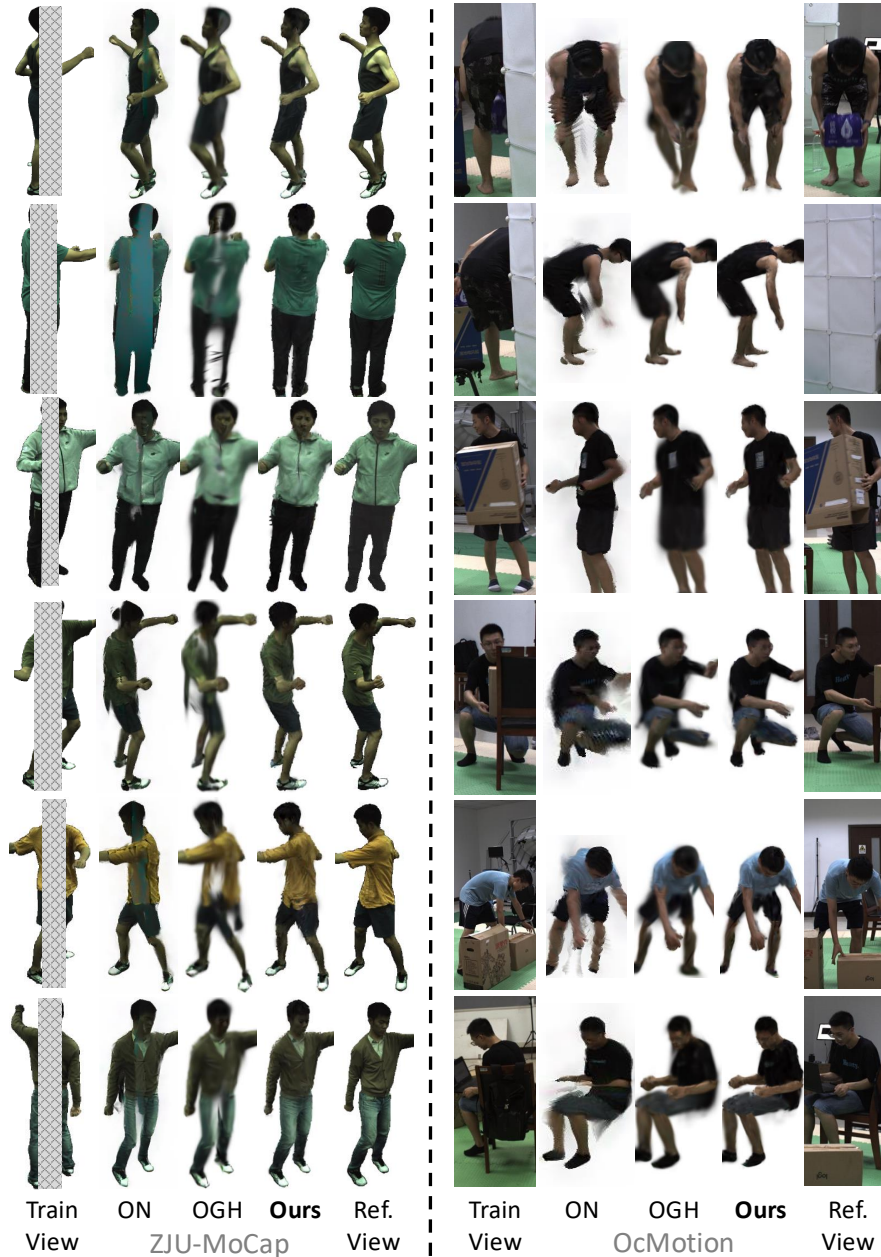


Figure 5: Qualitative comparisons on **simulated occlusions** in the ZJU-MoCap dataset [44] (left column) and **real-world occlusions** in the OcMotion dataset [16]. ON denotes OccNeRF [64] and OGH denotes OccGauHuman.

5.1 Datasets and Evaluation

ZJUMoCap. ZJU-MoCap [44] is a dataset consisting of 6 dynamic humans captured with a synchronized multi-camera system. Since the humans are in a lab environment free of occlusions, we follow OccNeRF’s [64] protocol to simulate occlusion of the human, masking out the center 50% of the human pixels for the first 80 % of frames. To challenge OccFusion on videos with even sparser frames, we use only **100 frames** from the first camera with a sampling rate of 5 to train the models and use the other 22 cameras for evaluation.

OcMotion. OcMotion [15] comprises of 48 videos of humans interacting with real objects in indoor environments. Experiments are conducted on the same 6 sequences adopted by Wild2Avatar [63], which are selected to provide a diverse coverage of real-world occlusions. We form sparser subsequences by sampling only 50 frames from each sequence to train the models.

Table 2: Ablation results on the ZJU-MoCap [44] dataset. LPIPS values are scaled by $\times 1000$.

Exp.	Methods	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Train time
-	GauHuman [14]	21.55	0.9430	55.88	10 mins
A	OccGauHuman	22.54	0.9457	54.88	2 mins
B	+ <i>Init Stage</i> generated masks $\{\hat{M}\}$	23.52	0.9516	52.35	5 mins
C	+ Posed space SDS	23.90	0.9510	55.47	7 mins
D	+ Canonical space SDS (<i>Optim Stage</i>)	23.91	0.9514	55.35	7 mins
E	+ <i>Refinement Stage</i>	23.96	0.9548	32.34	10 mins

Evaluation. We compare our OccFusion to OccNeRF [64], OccGaussian [67], and Wild2Avatar [63], the state-of-the-art in occluded human rendering. We also compare our results to GauHuman [14], HumanNeRF [57], and 3DGS-Avatar [46], popular human rendering methods not designed for occlusion. For fairness of comparison, all methods use the same set of segmentation masks and pose priors. We train GauHuman and OccGauHuman for 10 minutes each. We evaluate the methods both quantitatively and qualitatively. For our quantitative evaluations, we calculate the Peak Signal-to-Noise Ratio (PSNR), Structural SIMilarity (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS) metrics against the ground truth images. Since no ground truth is provided for OcMotion, we calculate the metrics on visible pixels only. For qualitative evaluations, we render the human from novel views and assess the quality of the renderings.

5.2 Results on Simulated and Real-world Occlusions

We provide quantitative metrics averaged over all the sequences in Table 1. Overall, methods designed for occluded human rendering tend to outperform their traditional counterparts. Among those methods, OccFusion consistently performs up to par or better than the state-of-the-art on both datasets while significantly beating all the baselines on LPIPS.

Qualitative results on novel view synthesis can be found in Figure 5. OccNeRF [64] has trouble generating unseen regions and renders significant discoloration and floaters when faced with occlusion. On the other hand, OccGauHuman’s renderings are blurry and occasionally incomplete. We observe that OccFusion is the only method to consistently render sharp and high-quality renderings free of occlusions.

5.3 Additional Studies

Ablation studies. We study the effect of each of our proposed components by adding them one by one and report average metrics on ZJU-MoCap in Table 2. Each stage plays a part towards optimal performance. Qualitative results on our ablations are included in Figure 6. We can see that the Initialization Stage helps enforce completeness for the initially incomplete human. The SDS regularization provided in the Optimization Stage helps remove floaters and artifacts in the posed and canonical space, further improving the shape of the human and enforcing completeness of the body. Finally, the Refinement Stage helps make the renderings more detailed in less observed regions, improving the rendering quality and greatly reducing the LPIPS.




Figure 6: Qualitative ablation studies. Please see Table 2 for corresponding experiments. Major differences are highlighted by red arrows.

Does the proposed OccGauHuman perform better than GauHuman [14] in rendering occluded humans? In section 3.3, we present a simple upgrade for the state-of-the-art 3DGS based human rendering model GauHuman [14] to help it better handle occlusions. Our improvements are straightforward but effective. We show quantitative results in Figure 1 (Left) and Table 1. As shown in Figure 7, our improved OccGauHuman reconstructs a more complete human body than the vanilla GauHuman.




Figure 7: Qualitative comparisons on **simulated occlusions** in the ZJU-MoCap dataset [44] (left column) and **real-world occlusions** in the OcMotion dataset [16] (right column). GH denotes GauHuman [14] and OGH denotes OccGauHuman.

6 Discussions and Conclusion

Limitations. Recovering occluded dynamic humans is challenging. As mentioned in section 4.3, reconstructing a 3D human requires adhering to multiple consistencies. However, even with the state-of-the-art generative models, it is still impossible to perfectly maintain those consistencies for 4D content (3D + motion) generation. Although our proposed methods are specifically designed to eliminate potential variances when using generative priors, we can still observe some generations are less coherent (e.g. Figure 4 and Figure 8), which may hurt the training of the rendering model on all stages. Moreover, we found that conditioning generative models with 2D poses is weak — the pose of the generated human does not always align with the condition pose, which may introduce even more uncertainty for training. In future work, we hope to train our own consistency-aware diffusion model specifically finetuned on human data.

Societal Impacts. Being able to reconstruct a human from an occluded monocular video can have a great societal impact. For example, having a high-fidelity 3D reconstruction of a human can help telemedicine practitioners become more immersed in the 3D space. While our research could lead to privacy concerns if humans are reconstructed without their consent, we believe that the benefits can be harnessed responsibly with appropriate safeguards.

Conclusion. In this work, we propose OccFusion, one of the first works that utilize 3D Gaussian splatting for occluded human rendering. Our approach consists of three stages: the Initialization, Optimization, and Refinement stages. By combining the efficiency and representative ability of 3D Gaussian splatting with the generation capabilities of diffusion priors, our method achieves state-of-the-art in occluded human rendering quality as measured by the PSNR, SSIM, and LPIPS metrics while only taking around 10 minutes to train. We hope our work inspires further exploration into the capabilities of diffusion priors to aid in human reconstruction.

7 Acknowledgment

This work was partially funded by the NIH Grant R01AG089169 and P41EB027060, Panasonic Holdings Corporation, the Gordon and Betty Moore Foundation, the Jaswa Innovator Award, Stanford HAI, Stanford HAI graduate fellowship, and Stanford Wu Tsai Human Performance Alliance.

References

- [1] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [2] Mingfei Chen, Jianfeng Zhang, Xiangyu Xu, Lijuan Liu, Yujun Cai, Jiashi Feng, and Shuicheng Yan. Geometry-guided progressive nerf for generalizable and efficient neural human rendering. In *European Conference on Computer Vision*, pages 222–239. Springer, 2022.
- [3] Xingyu Chen, Qi Zhang, Xiaoyu Li, Yue Chen, Ying Feng, Xuan Wang, and Jue Wang. Hallucinated neural radiance fields in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12943–12952, 2022.
- [4] Alvaro Collet, Ming Chuang, Pat Sweeney, Don Gillett, Dennis Evseev, David Calabrese, Hugues Hoppe, Adam Kirk, and Steve Sullivan. High-quality streamable free-viewpoint video. *ACM Transactions on Graphics (ToG)*, 34(4):1–13, 2015.
- [5] Mingsong Dou, Sameh Khamis, Yury Degtyarev, Philip Davidson, Sean Ryan Fanello, Adarsh Kowdle, Sergio Orts Escolano, Christoph Rhemann, David Kim, Jonathan Taylor, et al. Fusion4d: Real-time performance capture of challenging scenes. *ACM Transactions on Graphics (ToG)*, 35(4):1–13, 2016.
- [6] Chen Geng, Sida Peng, Zhen Xu, Hujun Bao, and Xiaowei Zhou. Learning neural volumetric representations of dynamic humans in minutes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8759–8770, 2023.
- [7] Rohit Girdhar, Mannat Singh, Andrew Brown, Quentin Duval, Samaneh Azadi, Sai Saketh Rambhatla, Akbar Shah, Xi Yin, Devi Parikh, and Ishan Misra. Emu video: Factorizing text-to-video generation by explicit image conditioning. *arXiv preprint arXiv:2311.10709*, 2023.
- [8] Shubham Goel, Georgios Pavlakos, Jathushan Rajasegaran, Angjoo Kanazawa, and Jitendra Malik. Humans in 4d: Reconstructing and tracking humans with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14783–14794, 2023.
- [9] Chen Guo, Tianjian Jiang, Xu Chen, Jie Song, and Otmar Hilliges. Vid2avatar: 3d avatar reconstruction from videos in the wild via self-supervised scene decomposition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12858–12868, 2023.
- [10] Kaiwen Guo, Peter Lincoln, Philip Davidson, Jay Busch, Xueming Yu, Matt Whalen, Geoff Harvey, Sergio Orts-Escolano, Rohit Pandey, Jason Dourgarian, et al. The relightables: Volumetric performance capture of humans with realistic relighting. *ACM Transactions on Graphics (ToG)*, 38(6):1–19, 2019.
- [11] Agrim Gupta, Lijun Yu, Kihyuk Sohn, Xiuye Gu, Meera Hahn, Li Fei-Fei, Irfan Essa, Lu Jiang, and José Lezama. Photorealistic video generation with diffusion models. *arXiv preprint arXiv:2312.06662*, 2023.
- [12] Liangxiao Hu, Hongwen Zhang, Yuxiang Zhang, Boyao Zhou, Boning Liu, Shengping Zhang, and Liqiang Nie. Gaussianavatar: Towards realistic human avatar modeling from a single video via animatable 3d gaussians. *arXiv preprint arXiv:2312.02134*, 2023.
- [13] Shoukang Hu, Fangzhou Hong, Liang Pan, Haiyi Mei, Lei Yang, and Ziwei Liu. Sherf: Generalizable human nerf from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9352–9364, 2023.
- [14] Shoukang Hu and Ziwei Liu. Gauhuman: Articulated gaussian splatting from monocular human videos. *arXiv preprint arXiv:2312.02973*, 2023.
- [15] Buzhen Huang, Yuan Shu, Jingyi Ju, and Yangang Wang. Occluded human body capture with self-supervised spatial-temporal motion prior. *arXiv preprint arXiv:2207.05375*, 2022.

- [16] Buzhen Huang, Tianshu Zhang, and Yangang Wang. Object-occluded human shape and pose estimation with probabilistic latent consistency. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [17] Boyi Jiang, Yang Hong, Hujun Bao, and Juyong Zhang. Selfrecon: Self reconstruction your digital avatar from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5605–5615, 2022.
- [18] Tianjian Jiang, Xu Chen, Jie Song, and Otmar Hilliges. Instantavatar: Learning avatars from monocular video in 60 seconds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16922–16932, 2023.
- [19] Wei Jiang, Kwang Moo Yi, Golnoosh Samei, Oncel Tuzel, and Anurag Ranjan. Neuman: Neural human radiance field from a single video. In *European Conference on Computer Vision*, pages 402–418. Springer, 2022.
- [20] HyunJun Jung, Nikolas Brasch, Jifei Song, Eduardo Perez-Pellitero, Yiren Zhou, Zhihao Li, Nassir Navab, and Benjamin Busam. Deformable 3d gaussian splatting for animatable human avatars. *arXiv preprint arXiv:2312.15059*, 2023.
- [21] Animesh Karnewar, Niloy J Mitra, Andrea Vedaldi, and David Novotny. Holofusion: Towards photo-realistic 3d generative modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22976–22985, 2023.
- [22] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. *arXiv preprint arXiv:2312.02145*, 2023.
- [23] Lei Ke, Mingqiao Ye, Martin Danelljan, Yu-Wing Tai, Chi-Keung Tang, Fisher Yu, et al. Segment anything in high quality. *Advances in Neural Information Processing Systems*, 36, 2024.
- [24] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4):1–14, 2023.
- [25] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023.
- [26] Muhammed Kocabas, Jen-Hao Rick Chang, James Gabriel, Oncel Tuzel, and Anurag Ranjan. Hugs: Human gaussian splats. *arXiv preprint arXiv:2311.17910*, 2023.
- [27] Youngjoong Kwon, Dahun Kim, Duygu Ceylan, and Henry Fuchs. Neural human performer: Learning generalizable radiance fields for human performance rendering. *Advances in Neural Information Processing Systems*, 34:24741–24752, 2021.
- [28] Mengtian Li, Shengxiang Yao, Zhifeng Xie, Keyu Chen, and Yu-Gang Jiang. Gaussianbody: Clothed human reconstruction via 3d gaussian splatting. *arXiv preprint arXiv:2401.09720*, 2024.
- [29] Mingwei Li, Jiachen Tao, Zongxin Yang, and Yi Yang. Human101: Training 100+ fps human gaussians in 100s from 1 view. *arXiv preprint arXiv:2312.15258*, 2023.
- [30] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 300–309, 2023.
- [31] Yuanze Lin, Ronald Clark, and Philip Torr. Dreampolisher: Towards high-quality text-to-3d generation via geometric diffusion. *arXiv preprint arXiv:2403.17237*, 2024.
- [32] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9298–9309, 2023.
- [33] Yang Liu, Xiang Huang, Minghan Qin, Qinwei Lin, and Haoqian Wang. Animatable 3d gaussian: Fast and high-quality reconstruction of multiple human avatars. *arXiv preprint arXiv:2311.16482*, 2023.
- [34] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. Syncdreamer: Generating multiview-consistent images from a single-view image. *arXiv preprint arXiv:2309.03453*, 2023.

- [35] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smp1: A skinned multi-person linear model. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 851–866. 2023.
- [36] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7210–7219, 2021.
- [37] 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. *Communications of the ACM*, 65(1):99–106, 2021.
- [38] Arthur Moreau, Jifei Song, Helisa Dharmo, Richard Shaw, Yiren Zhou, and Eduardo Pérez-Pellitero. Human gaussian splatting: Real-time rendering of animatable avatars. *arXiv preprint arXiv:2311.17113*, 2023.
- [39] Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, and Ying Shan. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 4296–4304, 2024.
- [40] 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.
- [41] Xiao Pan, Zongxin Yang, Jianxin Ma, Chang Zhou, and Yi Yang. Transhuman: A transformer-based human representation for generalizable neural human rendering. In *Proceedings of the IEEE/CVF International conference on computer vision*, pages 3544–3555, 2023.
- [42] Haokai Pang, Heming Zhu, Adam Kortylewski, Christian Theobalt, and Marc Habermann. Ash: Animatable gaussian splats for efficient and photoreal human rendering. *arXiv preprint arXiv:2312.05941*, 2023.
- [43] Bo Peng, Jun Hu, Jingtao Zhou, Xuan Gao, and Juyong Zhang. Intrinsicngp: Intrinsic coordinate based hash encoding for human nerf. *IEEE Transactions on Visualization and Computer Graphics*, 2023.
- [44] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9054–9063, 2021.
- [45] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022.
- [46] Zhiyin Qian, Shaofei Wang, Marko Mihajlovic, Andreas Geiger, and Siyu Tang. 3dgs-avatar: Animatable avatars via deformable 3d gaussian splatting. *arXiv preprint arXiv:2312.09228*, 2023.
- [47] Weining Ren, Zihan Zhu, Boyang Sun, Jiaqi Chen, Marc Pollefeys, and Songyou Peng. Nerf on-the-go: Exploiting uncertainty for distractor-free nerfs in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8931–8940, 2024.
- [48] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [49] Ruizhi Shao, Zerong Zheng, Hongwen Zhang, Jingxiang Sun, and Yebin Liu. Diffustereo: High quality human reconstruction via diffusion-based stereo using sparse cameras. In *European Conference on Computer Vision*, pages 702–720. Springer, 2022.
- [50] Zhuo Su, Lan Xu, Zerong Zheng, Tao Yu, Yebin Liu, and Lu Fang. Robustfusion: Human volumetric capture with data-driven visual cues using a rgbd camera. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, pages 246–264. Springer, 2020.
- [51] Guoxing Sun, Xin Chen, Yizhang Chen, Anqi Pang, Pei Lin, Yuheng Jiang, Lan Xu, Jingyi Yu, and Jingya Wang. Neural free-viewpoint performance rendering under complex human-object interactions. In *Proceedings of the 29th ACM International Conference on Multimedia, MM '21*, page 4651–4660, New York, NY, USA, 2021. Association for Computing Machinery.
- [52] Wenzhang Sun, Yunlong Che, Han Huang, and Yandong Guo. Neural reconstruction of relightable human model from monocular video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 397–407, 2023.

- [53] 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.
- [54] Guangcong Wang, Zhaoxi Chen, Chen Change Loy, and Ziwei Liu. Sparsenerf: Distilling depth ranking for few-shot novel view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9065–9076, 2023.
- [55] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12619–12629, 2023.
- [56] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [57] Chung-Yi Weng, Brian Curless, Pratul P Srinivasan, Jonathan T Barron, and Ira Kemelmacher-Shlizerman. Humannerf: Free-viewpoint rendering of moving people from monocular video. In *Proceedings of the IEEE/CVF conference on computer vision and pattern Recognition*, pages 16210–16220, 2022.
- [58] Zhenzhen Weng, Jingyuan Liu, Hao Tan, Zhan Xu, Yang Zhou, Serena Yeung-Levy, and Jimei Yang. Single-view 3d human digitalization with large reconstruction models. *arXiv preprint arXiv:2401.12175*, 2024.
- [59] Zhenzhen Weng, Zeyu Wang, and Serena Yeung. Zeroavatar: Zero-shot 3d avatar generation from a single image. *arXiv preprint arXiv:2305.16411*, 2023.
- [60] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7623–7633, 2023.
- [61] Rundi Wu, Ben Mildenhall, Philipp Henzler, Keunhong Park, Ruiqi Gao, Daniel Watson, Pratul P Srinivasan, Dor Verbin, Jonathan T Barron, Ben Poole, et al. Reconfusion: 3d reconstruction with diffusion priors. *arXiv preprint arXiv:2312.02981*, 2023.
- [62] Jamie Wynn and Daniyar Turmukhambetov. Diffusionnerf: Regularizing neural radiance fields with denoising diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4180–4189, 2023.
- [63] Tiange Xiang, Adam Sun, Scott Delp, Kazuki Kozuka, Li Fei-Fei, and Ehsan Adeli. Wild2avatar: Rendering humans behind occlusions. *arXiv preprint arXiv:2401.00431*, 2023.
- [64] Tiange Xiang, Adam Sun, Jiajun Wu, Ehsan Adeli, and Li Fei-Fei. Rendering humans from object-occluded monocular videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3239–3250, 2023.
- [65] Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. *arXiv preprint arXiv:2404.07191*, 2024.
- [66] Xiaofeng Yang, Yiwen Chen, Cheng Chen, Chi Zhang, Yi Xu, Xulei Yang, Fayao Liu, and Guosheng Lin. Learn to optimize denoising scores for 3d generation: A unified and improved diffusion prior on nerf and 3d gaussian splatting. *arXiv preprint arXiv:2312.04820*, 2023.
- [67] Jingrui Ye, Zongkai Zhang, Yujiao Jiang, Qingmin Liao, Wenming Yang, and Zongqing Lu. Occgaussian: 3d gaussian splatting for occluded human rendering. *arXiv preprint arXiv:2404.08449*, 2024.
- [68] Keyang Ye, Tianjia Shao, and Kun Zhou. Animatable 3d gaussians for high-fidelity synthesis of human motions. *arXiv preprint arXiv:2311.13404*, 2023.
- [69] Vickie Ye, Georgios Pavlakos, Jitendra Malik, and Angjoo Kanazawa. Decoupling human and camera motion from videos in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 21222–21232, 2023.
- [70] Taoran Yi, Jiemin Fang, Junjie Wang, Guanjun Wu, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Qi Tian, and Xinggang Wang. Gaussiandreamer: Fast generation from text to 3d gaussians by bridging 2d and 3d diffusion models, 2024.
- [71] Tao Yu, Zerong Zheng, Kaiwen Guo, Pengpeng Liu, Qionghai Dai, and Yebin Liu. Function4d: Real-time human volumetric capture from very sparse consumer rgb-d sensors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5746–5756, 2021.

- [72] Zhengming Yu, Wei Cheng, Xian Liu, Wayne Wu, and Kwan-Yee Lin. Monohuman: Animatable human neural field from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16943–16953, 2023.
- [73] Zhongrui Yu, Haoran Wang, Jinze Yang, Hanzhang Wang, Zeke Xie, Yunfeng Cai, Jiale Cao, Zhong Ji, and Mingming Sun. Sgd: Street view synthesis with gaussian splatting and diffusion prior. *arXiv preprint arXiv:2403.20079*, 2024.
- [74] Ye Yuan, Xueting Li, Yangyi Huang, Shalini De Mello, Koki Nagano, Jan Kautz, and Umar Iqbal. Gavatar: Animatable 3d gaussian avatars with implicit mesh learning. *arXiv preprint arXiv:2312.11461*, 2023.
- [75] Dongbin Zhang, Chuming Wang, Weitao Wang, Peihao Li, Minghan Qin, and Haoqian Wang. Gaussian in the wild: 3d gaussian splatting for unconstrained image collections. *arXiv preprint arXiv:2403.15704*, 2024.
- [76] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023.
- [77] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [78] Tingyang Zhang, Qingzhe Gao, Weiyu Li, Libin Liu, and Baoquan Chen. Bags: Building animatable gaussian splatting from a monocular video with diffusion priors. *arXiv preprint arXiv:2403.11427*, 2024.
- [79] Zhizhuo Zhou and Shubham Tulsiani. Sparsefusion: Distilling view-conditioned diffusion for 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12588–12597, 2023.

A Table of Symbols

For notation simplicity, we adopted alphabetic symbols in this paper to represent essential components in our framework. For better symbol-name correspondences, here we justify the implications of all symbols used in the paper in Table 3.

Table 3: Table of symbols.

Symbols	Explanations
<i>Preliminaries</i>	
\mathbf{x}_c	3D points in the canonical human space
\mathbf{x}_P	3D points in the posed human space
w	skinning weights used in LBS
G	transformation matrix used in LBS
b	translation vector used in LBS
\mathbf{J}	3D locations of human joints
θ	pose parameters used in SMPL [35]
β	shape parameters used in SMPL [35]
\mathbf{p}	center of a 3D Gaussian
o	opacity of a 3D Gaussian
μ	mean value of a 3D Gaussian
Σ	covariance matrix of a 3D Gaussian
<i>OccFusion</i>	
Π	optimizable human 3D Gaussians
Φ	a pretrained generative model [48], used as prior
\mathbf{M}	precomputed binary human mask, used as prior
\mathbf{P}	precomputed human pose, used as prior
$\hat{\mathbf{P}}$	the canonical articulation of \mathbf{P}
\mathbf{I}	input image with occluded human
$\hat{\mathbf{M}}$	Init Stage generated complete human mask
Δ	SDS gradients, used as a guidance in the Optim Stage
$\hat{\mathbf{I}}$	Optim Stage rendered human RGB image
\mathbf{A}	Π rendered human occupancy map in all stages
\mathbf{C}	Refine Stage rendered human RGB image
\mathbf{R}	inpainting mask computed by $(1 - \mathbf{M}) \cdot \mathbf{A}$
ρ	a random variable $\in [0, 1]$ controls Optim Stage SDS

B Implementation Details

OccFusion requires several priors. We run SAM [25] to get all the human masks $\{\mathbf{M}\}$. While we follow previous work [64, 63] and use the ground truth poses provided by ZJU-MoCap and OcMotion, pose priors \mathbf{P} can be obtained via occlusion-robust SMPL prediction/optimization methods such as HMR 2.0 [8] and SLAHMR [69] for in-the-wild videos. Improving the quality of priors is not the focus of this work. We use the pre-trained Stable Diffusion 1.5 model [48] with ControlNet [76] plugins for SDS in all the stages.

In the *Initialization Stage*, instead of inpainting incomplete human masks directly, we run the pretrained diffusion model to inpaint RGB images with 10 inference steps and 1.0 ControlNet conditioning scale. We use the positive prompt — "*clean background, high contrast to the background, a person only, plain clothes, simple clothes, natural body, natural limbs, no texts, no overlay*" and the negative prompt — "*multiple objects, occlusions, complex pattern, fancy clothes, longbody, lowres, bad anatomy, bad hands, bad feet, missing fingers, cropped, worst quality, low quality, blurry*". After inpainting the RGB images, we then run SAM-HQ [23] with \mathbf{P} as the prompts to get $\{\hat{\mathbf{M}}\}$.

In the *Optimization Stage*, we train the 3D human Gaussian Π from scratch by following the objective Equation 5. We set $\lambda_{rgb} = 1e^4$, $\lambda_{mask} = 2e^4$, $\lambda_{ssim} = 1e^3$, and $\lambda_{lips} = 1e^3$. At each training step, we random switch the SDS regularization on either posed human space or the canonical Da-pose space with a probability of 75% and 25%. When applying SDS regularization on the canonical human

space, we randomly rotate the human horizontally with a uniformly sampled degree in $\{k\frac{\pi}{9}, k \in \mathbb{Z}\}$. We set the SDS loss weights as $\lambda_{pose} = 2e^5$ and $\lambda_{can} = 2e^5$. In this stage, we train Π for 1200 steps.

In the *Refinement Stage*, we first generate the RGB human inpaintings via the proposed in-context inpainting method. We run the pretrained diffusion model with conditions on \mathbf{M} , 10 inference steps, and 0.3 ControlNet conditioning scale. We did not use positive prompts for the inpainting but used the same negative prompts as in the *Optimization Stage*. During training, we set the loss weights as $\lambda_{rgb} = 1$ and $\lambda_{mask} = 0.1$, $\lambda_{gen} = 0.1$, and $\lambda_{lips} = 0.2$. In this stage, we finetune Π for another 1800 steps with Gaussian densification and pruning enabled for the first 1000 steps.

C Additional Studies

Effectiveness of in-context inpainting. We provide comparisons of the human in the Refinement Stage with and without in-context inpainting and provide qualitative comparisons in Figure 8. While renderings from the Optimization stage are less detailed in occluded areas, our proposed in-context inpainting is able to generate the missing content and greatly increase the rendering quality in these areas.



Figure 8: Comparison of the inpainted human in the Refinement Stage with and without using the proposed in-context inpainting technique. Major differences are highlighted with red arrows.




Figure 9: Applying SDS on RGB images vs. on human occupancy maps. As mentioned in Sec. 4.1 and Fig. 4 of the main paper, generated RGB appearances are much more inconsistent than generated silhouettes. As a result, applying SDS on RGB leads to defective rendering results.




Figure 10: Training with complete unoccluded masks vs. with inpainted masks in the Optim. stage. Although inpainted masks are slightly more inconsistent compared to the complete masks, our training pipeline converges to the same level of rendering quality.




Figure 11: Novel view synthesis results from InstantMesh [65] conditioned on the least occluded frame. Discrepancies are circled in red.

Applying SDS on RGB vs on Human Occupancy Maps. We include additional experiments comparing the rendering results of applying SDS on RGB vs. on human occupancy maps (as proposed). It is clear that applying SDS on RGB leads to defective renderings as well as inferior quantitative results. This experiment validates our claim made in Sec. 4.1 and Fig. 4 in the main paper.

Robustness of training to inpainted masks. For in-the-wild occluded videos, there are no ground truth masks for the occluded body regions due to unknown human/garment deformations. Relying on the state-of-the-art pre-trained priors brought by the Segment Anything model (SAM) [25] and Stable Diffusion [48], the segmented/inpainted masks are expected to be reasonable and coherent across frames. To test the robustness of our method to variances in the in-painted masks, we add comparison experiments on ZJU-MoCap that supervise using the complete SAM masks obtained from the unoccluded humans with minimum variances. Please see the qualitative results in Figure 10. We find that using the inpainted \hat{M} leads to a good enough rendering quality comparable to using masks derived from the unoccluded images, validating the robustness of our model.

Can existing generative models recover an occluded human? While there are works for using generative diffusion models to render 3D humans conditioned on single [58] and multiple [49] images, none are able to condition on a monocular video of the person.

Since [58] has not released code, we include results from InstantMesh [65]. We use the provided segmentation mask to mask the least occluded frame onto a white background and use it as conditioning. Novel view synthesis results are included in Figure 11. InstantMesh is unable to recover a complete human geometry and fails to generate a reasonable appearance from the single image.

D Video Studies

For a more comprehensive presentation of the results, we include video renderings on all the training frames for both datasets. For the ZJU-MoCap videos (named with the prefix `zju`), from left to right, we show the occluded human, OccGauHuman rendering, Optimization Stage rendering, Refinement Stage rendering, and the reference. For the OcMotion videos (named with the prefix `ocmotion`), without references for real-world occlusions, from left to right, we show the occluded human, OccGauHuman rendering, Optimization Stage rendering, and Refinement Stage rendering.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Our work proposes a novel three-stage pipeline for rendering occluded humans with Gaussian splatting and diffusion priors that outperforms the state-of-the-art in the field.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Our paper includes a "Limitations" section in Section 6 that acknowledges the inconsistency of generative diffusion models, which are central to our approach.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: Our paper does not include any new theorems, formulas, or proofs.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: ZJU-MoCap and OcMotion, the datasets used in this paper, are publicly available, and we include all implementation details in supplementary materials.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We are not including code in our submission. However, we provide implementation details in our supplementary materials.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Implementation details are all included in supplementary materials.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Error bars are not reported, since they are not a common practice in the human rendering research community.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We mention the hardware used to train our method at the end of Section 4.3. We also break down the amount of time needed to train each stage in Table 2.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The authors have reviewed the NeurIPS Code of Ethics and strongly believe it conforms to the Code.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We include discussion of societal impacts in 6.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our method utilizes per-video optimization, so no pretrained weights are provided and no models or data are released.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: In Subsec 3.3, we cite GauHuman in our paper and the license used, and provide a link to the code.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: New assets are not used or included in our submission.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our work does not involve crowdsourcing or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Our work does not involve crowdsourcing or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.