

Deep Inverse Shading: Consistent Albedo and Surface Detail Recovery via Generative Refinement

Jiacheng Wu, Ruiqi Zhang, Jie Chen[†]

Department of Computer Science, Hong Kong Baptist University, Hong Kong SAR, China
 csjcwu@comp.hkbu.edu.hk, csrqzhang@comp.hkbu.edu.hk, chenjie@comp.hkbu.edu.hk

Abstract

Reconstructing human avatars using generative priors is essential for achieving versatile and realistic avatar models. Traditional approaches often rely on volumetric representations guided by generative models, but these methods require extensive volumetric rendering queries, leading to slow training. Alternatively, surface-based representations offer faster optimization through differentiable rasterization, yet they are typically limited by vertex count, restricting mesh resolution and scalability when combined with generative priors. Moreover, integrating generative priors into physically based human avatar modeling remains largely unexplored. To address these challenges, we introduce DIS (Deep Inverse Shading), a unified framework for high-fidelity, relightable avatar reconstruction that incorporates generative priors into a coherent surface representation. DIS centers on a mesh-based model that serves as the target for optimizing both surface and material details. The framework fuses multi-view 2D generative surface normal predictions, rich in detail but often inconsistent, into the central mesh using a normal conversion module. This module converts generative normal outputs into per-triangle surface offsets via differentiable rasterization, enabling the capture of fine geometric details beyond sparse vertex limitations. Additionally, DIS integrates a de-shading module, informed by generative priors, to recover accurate material properties such as albedo. This module refines albedo predictions by removing baked-in shading and back-propagates reconstruction errors to further optimize the mesh geometry. Through this joint optimization of geometry and material appearance, DIS achieves physically consistent, high-quality reconstructions suitable for accurate relighting. Our experiments show that DIS delivers SOTA relighting quality, enhanced rendering efficiency, lower memory consumption, and detailed surface reconstruction.

Extended version — <https://arxiv.org/abs/2511.08079>

Introduction

Reconstructing human avatars is essential for immersive applications such as 3D movies and games. Traditional methods focus on recovering geometry from captured data, while recent advances in generative models enable the integration

of learned priors, enhancing detail and consistency. This hybrid approach offers a promising direction for building universal, relightable avatars.

A straightforward approach that builds on generative models involves synthesizing multi-view images and subsequently fusing them into 3D geometry (Long et al. 2023; Voleti et al. 2024; Li et al. 2025), extending traditional reconstruction pipelines by incorporating learned priors to enrich view synthesis. However, a key limitation of this pipeline is the lack of 3D consistency in the generated views, which leads to challenges in accurately recovering high-fidelity geometries, particularly for detailed human avatars.

Building on this observation, recent methods employ a central 3D volumetric representation guided by 2D generative models. This strategy leverages the rich priors of generative models while maintaining spatial coherence for more accurate and consistent reconstructions. Score Distillation Sampling (SDS) loss (Poole et al. 2022; Huang et al. 2024a,b) is widely used in this paradigm, refining volumetric fields with supervision from 2D diffusion models. However, volumetric rendering introduces significant computational overhead and slow training, prompting a growing interest in more efficient alternatives.

When it comes to geometry representation for avatar reconstruction, volumetric approaches remain prevalent. Neural implicit fields (Mildenhall et al. 2020), for example, encode geometry using coordinate-based neural networks, enabling continuous surface representation but requiring dense sampling and incurring high computational costs. Recently, 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) has improved rendering speed by approximating surfaces using sparse Gaussian primitives. However, it sacrifices geometric precision, which is critical for accurate visibility computation and physically-based rendering (PBR).

To address the computational inefficiency and geometric imprecision of volumetric methods, surface-based representations such as meshes have emerged as promising alternatives for PBR tasks. These representations offer explicit surface definitions, enabling efficient rendering and accurate geometric modeling. Parametric human models (e.g., SMPL (Loper et al. 2015) and SMPL-X (Pavlakos et al. 2019)) offer a deformable mesh structure suitable for animation, but their coarse topology limits capture of high-frequency details essential for photorealistic appearance.

[†]Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Consequently, recent efforts focus on mesh-based strategies that balance computational efficiency with reconstruction fidelity, though challenges remain in optimizing detailed surface variations over sparse mesh topologies.

To address the challenge of modeling fine-grained surface detail on sparse mesh topologies, we propose a surface offset network, inspired by micro-mesh representations (Maggiordomo, Moreton, and Tarini 2023; Dou et al. 2024), to capture subtle geometric variations across a dynamic mesh. Our core design is a normal conversion module that maps UV-space surface offsets into 3D normal vectors. This enables gradients from multi-view 2D normal predictions, generated by generative models, to be propagated directly onto the 3D mesh via differentiable rasterization, thereby unifying multi-view signals onto a coherent and optimizable 3D surface. Notably, we represent surface details as surface variations that do not alter the underlying mesh geometry.

While the integration of 2D generative normal introduces valuable fine-grained detail, these predictions often lack angular consistency and exhibit temporal instability. To address this, our method explicitly optimizes the mesh surface against these noisy yet informative generations, enabling the construction of a coherent and high-resolution surface representation. Compared to volumetric approaches, our surface-centric framework provides a more computationally efficient and geometrically precise central representation for physically-based rendering applications.

Nevertheless, integrating generative priors into surface material modeling for PBR in human avatar modeling remains a significant challenge. Existing methods (Kim et al. 2024; He et al. 2024) for general objects often rely on large-scale datasets, but struggle to disentangle shading from intrinsic reflectance properties. To address this, we introduce a prior-driven de-shading module that estimates albedo by removing baked-in shading, leveraging generative shading priors conditioned on geometry and texture. The refined albedo enables reconstruction errors to be back-propagated, allowing for the joint optimization of geometry and appearance. This results in a co-optimization loop that maintains physical consistency across both shape and material.

The contributions are summarized as follows:

- We propose a mesh-based central surface representation that efficiently reconstructs geometry from multi-view generations. A normal conversion module transforms 2D normal predictions into UV-space surface offsets, enabling gradient flow from 2D to 3D via differentiable rasterization, achieving finer detail than volumetric representations with significantly lower computational cost.
- We introduce a prior-driven de-shading module that recovers albedo by removing shading, using priors conditioned on surface texture and geometry. This facilitates material disentanglement and captures fine-grained surface detail for physically plausible rendering.
- We present a unified co-optimization strategy that jointly refines surface geometry and intrinsic appearance. By aligning multi-view consistency with generative priors and coupling geometry with shading in a physically-based rendering loop, our method achieves SOTA relight-

ing accuracy, efficiency, and surface fidelity.

Related Works

Volumetric & Surface Representation

Volumetric representations (Mildenhall et al. 2020; Wang et al. 2021; Müller et al. 2022) model the whole scene and are widely applied in human avatar modeling frameworks. Classical implicit neural fields utilize coordinate-based MLPs to model dynamic human avatars (Gafni et al. 2021; Pumarola et al. 2021) by mapping spatial coordinates into implicit neural fields. These methods often combine implicit neural fields with parametric human models to leverage surface constraints for dynamic human avatar modeling (Wang et al. 2022; Zhang and Chen 2022; Peng et al. 2024b). However, the extensive querying required for implicit fields leads to long training times and complicates efficient PBR. While 3DGS is efficient to manipulate and apply in human avatar modeling (Qian et al. 2024; Zhang and Chen 2024), it encounters memory challenges as resolution increases, and its sparse nature makes it hard to build a precise surface necessary for accurate PBR calculation. In contrast, surface-based approaches focus on capturing surface variations and are commonly used in dynamic human avatar modeling. These methods typically employ parametric models like SMPL (Loper et al. 2015) and SMPL-X (Pavlakos et al. 2019) to improve memory efficiency (Zhang, Chen, and Wang 2023). However, the limited number of vertices in these models restricts their ability to represent fine surface details, which is crucial for high-quality PBR. Recent advances introduce micro-mesh representations (Maggiordomo, Moreton, and Tarini 2023; Dou et al. 2024), which add displacement to low-polygon meshes as micro-vertices, improving surface detail without significant memory cost. However, effectively applying micro-meshes to dynamic human avatars remains underexplored. To bridge this gap, we use an offset network to recover micro-vertex displacement, represented as surface variations for dynamic human avatars. We also introduce a normal conversion module that links 2D image features to the 3D mesh surface, resulting in a deformed parametric model with intricate surface variations to be incorporated into PBR for realistic rendering.

3D Human Reconstruction with Generative Priors

Diffusion models (Rombach et al. 2022) have demonstrated significant potential in image generation (Hu et al. 2023; Zhu et al. 2024). Recent research has investigated their application in multi-view human images, with the goal of fusing these images for avatar geometry reconstruction. While some methods attempt to tackle the multi-view inconsistency by integrating 3D information into the 2D diffusion process (Long et al. 2023; Voleti et al. 2024; Li et al. 2025), they still struggle with multi-view inconsistencies. This limitation arises because these methods primarily focus on enhancing 2D image consistency rather than addressing the underlying 3D object coherence. Thus, a more effective strategy is to utilize a base 3D representation and optimize it with priors from 2D diffusion models. This approach promises more accurate and consistent 3D reconstructions for creating

high-quality human avatars. Many current methods employ a pre-trained 2D diffusion model to guide volumetric-centric human representation via SDS loss (Huang et al. 2024b; Xiu et al. 2024). However, the extensive sampling point queries slow down training, and the SDS’s lack of 3D structure awareness may lead to instability, exemplified by issues like the multi-face Janus problem (Poole et al. 2022). To address these challenges, our primary goal is to effectively leverage information from prior-driven models to improve surface-centric representation while stably ensuring robust underlying geometry during human avatar reconstruction. Specifically, we introduce a normal conversion module that transforms surface variations into normals, allowing us to efficiently distill normal enhancement priors from a pre-trained diffusion model to refine coarse normals during reconstruction. Moreover, we employ inverse shading priors from a prior-driven de-shading module to rectify geometry and material during material learning. By effectively integrating prior-driven models with human avatar reconstruction, our method accurately models mesh surfaces, disentangles precise material information, and produces more realistic relighting results.

Relightable Human Avatar Reconstruction

Illuminating objects and human avatars involves acquiring novel lighting settings and understanding material properties (Shu et al. 2017; Guo et al. 2019). Traditional methods (Debevec et al. 2000; Habermann et al. 2019) estimate these properties using photometric stereo techniques under controlled lighting conditions. However, complex real-world illumination poses challenges to these approaches. To address this limitation, recent advancements have shifted towards learning-based techniques that leverage neural rendering (Li et al. 2024; Chen et al. 2024) and neural inverse rendering (Srinivasan et al. 2021; Zhang et al. 2021; Jin et al. 2023) to create relightable models grounded in geometric assumptions. These methods inspire research that decomposes human figures into canonical reflectance fields and incorporates human templates for motion modeling (Chen and Liu 2022; Sun et al. 2023; Iqbal et al. 2023; Wu et al. 2025). Despite their potential, the lack of geometric supervision often leads to inaccuracies in surface reconstruction. More discussions are presented in the extended version. To improve efficiency, our framework facilitates material disentanglement using an explicit mesh with intricate surface variations. By rectifying geometry and material information guided by a prior-driven model, we achieve precise surface information querying and accurate material disentanglement.

The Proposed Method

We propose Deep Inverse Shading (DIS), a powerful framework for reconstructing high-fidelity human avatars from sparse-view videos. DIS jointly learns detailed geometry and material properties using generative priors within a unified, surface-based representation. Built on a dynamic SMPL mesh (Loper et al. 2015), DIS refines geometry via a normal conversion module that maps 2D normal predictions to 3D surfaces, and disentangles shading from albedo with

a generative de-shading module for physically plausible appearance modeling, using multi-view generative predictions.

The rest of this section is organized as follows. First, we present our prior-driven mesh optimization approach via a normal conversion module, which integrates 2D generative normal priors into 3D meshes using differentiable rasterization. Next, we introduce our physically-based rendering strategy via differentiable surface integration, detailing how material properties are combined with geometry for realistic shading. Finally, we explain our joint material-geometry refinement method through generative inverse shading, which leverages learned priors to co-optimize geometry and appearance in a physically consistent manner.

Prior-Driven Mesh Optimization via Generative Normal Conversion

Differentiable rasterization (Ravi et al. 2020) enables gradients to flow from 2D image space to the 3D mesh, allowing effective supervision of 3D geometry from 2D observations. However, traditional parametric models have sparse vertices, limiting their ability to capture fine details. To overcome this, we initialize a coarse mesh G_{coarse} and introduce a normal conversion module $O2N$, which refines triangle-level surface normals using an offset network. These offsets are defined in UV space and guided by a prior-driven model conditioned on RGB images and coarse normal predictions.

We first fit an SMPL model to each frame of the input sparse-view videos to obtain a base mesh for deformation. An offset network $\mathcal{M}_{\text{offset}}$ then predicts scalar offsets l_v for each SMPL vertex $\hat{\mathbf{x}}_v$, applied along vertex normals \mathbf{n}_v to deform the mesh into G_{coarse} , a coarse approximation of the human surface: $\hat{\mathbf{x}}_v = \mathbf{x}_v + \mathbf{n}_v \cdot l_v$.

Using a differentiable rasterizer, we establish pixel-wise correspondences between image-space coordinates (x, y) and texel-space coordinates (u, v) on the mesh under a given pose θ . This mapping enables efficient access to canonical surface attributes at each on-screen pixel. Specifically, for each mesh surface point, we extract interpolated rasterized values such as normal vector $\mathbf{n} \in \mathbb{R}^3$, 3D location \mathbf{x} , and query the scalar offset value l from the same offset network $\mathcal{M}_{\text{offset}}$. In this way, we align the UV coordinates of on-screen pixels, SMPL vertices, and intra-triangle points within a unified UV space. This setup allows us to deform not only the SMPL vertices but also the sub-vertex points inside each triangle to compute 3D surface positions \mathbf{x}_{surf} :

$$\mathbf{x}_{\text{surf}} = \mathbf{x} + \mathbf{n} \cdot l. \quad (1)$$

These positions are computed *per-pixel* on the rasterizer for pixel-wise detailed geometric refinement.

The process of normal conversion is illustrated in Fig. 2, where a representative 2D UV-space patch of five rasterized pixel points—denoted as m, i, j, k, l is used to estimate the surface normal at the center point i . Their offset values are converted into a coherent normal vector using Eq.(2), capturing the local surface direction variation:

$$\begin{aligned} \mathbf{n}_i &= \mathbf{n}_{\Delta_{ijk}} + \mathbf{n}_{\Delta_{ikl}} + \mathbf{n}_{\Delta_{ilm}} + \mathbf{n}_{\Delta_{imj}} \\ &= (\mathbf{x}_j - \mathbf{x}_i) \times (\mathbf{x}_k - \mathbf{x}_i) + (\mathbf{x}_k - \mathbf{x}_i) \times (\mathbf{x}_l - \mathbf{x}_i) \\ &\quad + (\mathbf{x}_l - \mathbf{x}_i) \times (\mathbf{x}_m - \mathbf{x}_i) + (\mathbf{x}_m - \mathbf{x}_i) \times (\mathbf{x}_j - \mathbf{x}_i), \\ \bar{\mathbf{n}}_i &= \frac{\mathbf{n}_i}{\|\mathbf{n}_i\|_2}. \end{aligned} \quad (2)$$

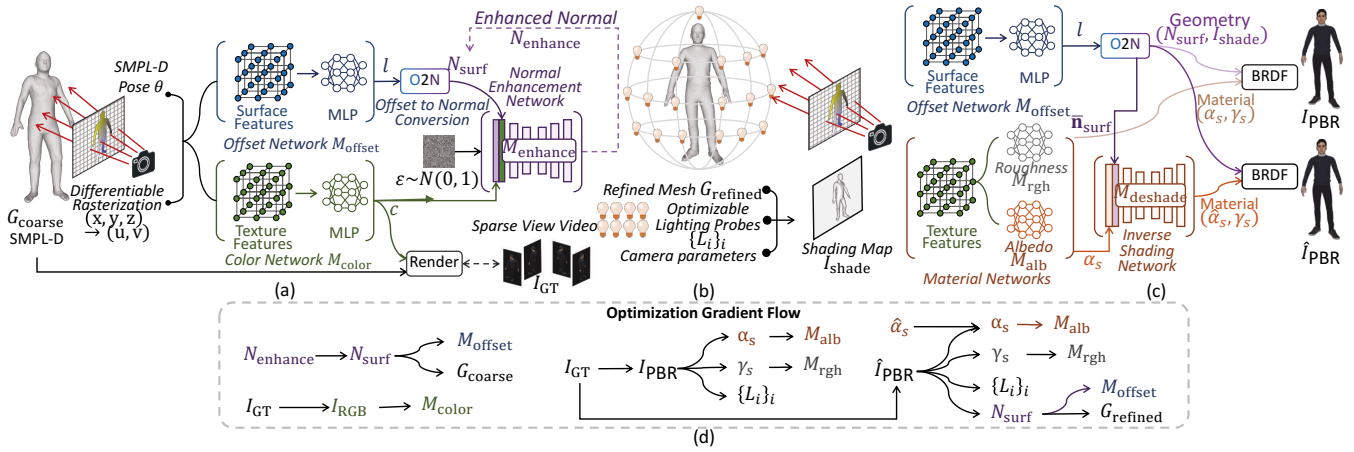


Figure 1: System diagram: (a) Prior-driven mesh optimization. The pipeline starts by estimating vertex offsets with the surface offset network $\mathcal{M}_{\text{offset}}$ to deform a coarse mesh G_{coarse} . Differentiable rasterization provides pixel UV coordinates, which are converted to surface normals N_{surf} via the normal conversion module $O2N$. The color network $\mathcal{M}_{\text{color}}$ predicts pixel color c and image I_{RGB} , while the normal enhancement model $\mathcal{M}_{\text{enhance}}$ refines N_{surf} to enhanced normals N_{enhance} , producing a detailed dynamic mesh G_{refined} . (b) Surface-based physically-based rendering. Learnable light probes $\{L_i\}_i$ are placed around G_{refined} to estimate visibility. With rasterized surface attributes, we compute albedo α_s via \mathcal{M}_{alb} , roughness γ_s via \mathcal{M}_{rgh} (both conditioned on shared features from $\mathcal{M}_{\text{color}}$), and N_{surf} . These are used in BRDF-based rendering to synthesize the image I_{PBR} . (c) Inverse shading and refinement. To remove baked-in lighting from α_s , a de-shading module estimates a clean albedo $\hat{\alpha}_s$ conditioned on α_s and N_{surf} , and computes an improved relit image \hat{I}_{PBR} . (d) Joint optimization. Gradients are propagated through the entire pipeline: (1) $\mathcal{M}_{\text{color}}$ is supervised by I_{RGB} ; (2) $\mathcal{M}_{\text{offset}}$ and G_{coarse} by N_{enhance} ; (3) \mathcal{M}_{alb} , \mathcal{M}_{rgh} , and $\{L_i\}_i$ by I_{PBR} ; (4) the same modules plus G_{refined} by \hat{I}_{PBR} ; and (5) \mathcal{M}_{alb} is further refined using $\hat{\alpha}_s$.




Figure 2: Normal Conversion Module. This module first queries rasterized pixel-level information, including pixel UV coordinates, normals, and world coordinates. Pixel offsets, predicted by an offset MLP, are integrated with the pixel data to compute 3D surface coordinates via Eq. (1), capturing fine-grained surface variations within each triangle. Finally, point normals are calculated by constructing four surrounding triangles, as described in Eq. (2).

To obtain the final surface normal map N_{surf} , we integrate the per-pixel surface normal \bar{n}_{surf} . In parallel, we query a dynamic color network $\mathcal{M}_{\text{color}}$ to compute the pixel color c of the coarse mesh G_{coarse} under the current pose θ at each texel coordinate (u, v) , integrated as the RGB image I_{RGB} . $\mathcal{M}_{\text{offset}}$ and $\mathcal{M}_{\text{color}}$ each consist of a hash grid encoding followed by a 5-layer MLP, as in (Müller et al. 2022).

To improve both the coarse geometry G_{coarse} and the normal map N_{surf} , we introduce a normal enhancement model $\mathcal{M}_{\text{enhance}} : (N_{\text{surf}}, I_{\text{RGB}}) \rightarrow N_{\text{enhance}}$ that predicts detailed nor-

mals from the surface normal image N_{surf} and color image I_{RGB} . This model provides prior-driven supervision beyond the limited RGB information from sparse-view video. Efficiently leveraging the graphics pipeline, we perform parallel texel-wise queries for all on-screen pixels. The resulting enhanced normals N_{enhance} are used to refine N_{surf} , and, via the normal conversion module, improvements are back-propagated to optimize both the surface normal information in $\mathcal{M}_{\text{offset}}$, and vertex accuracy of G_{coarse} to obtain G_{refined} .

To implement $\mathcal{M}_{\text{enhance}}$, we adopt Stable Diffusion v1.5 (Rombach et al. 2022) as the backbone, pre-trained on a synthetic dataset to generate enhanced normals N_{enhance} . A modified ControlNeXt (Peng et al. 2024a) fuses coarse normals and RGB images as conditional inputs, enabling detailed pixel-level refinement. Rather than generating normals from scratch, the model enhances low-resolution predictions. To handle domain shifts and preserve learned knowledge, only a subset of weights within Stable Diffusion is fine-tuned. The training data, curated with help from a vision foundation model (Khrodkar et al. 2024), includes coarse normals, RGB images, and high-quality ground truth normals.

Physically-Based Rendering via Differentiable Surface Integration

Accurately modeling a physically relightable human avatar requires consideration of illumination, geometry, and surface material properties. Central to capturing their interaction is the Bidirectional Reflectance Distribution Function (BRDF) (Walter et al. 2007), which describes how light re-

flects off opaque surfaces. The BRDF defines the reflection of incoming light at various angles based on surface attributes, viewing direction, and material properties. The physically-based rendering process is defined in Eq. (3):

$$I_{\text{PBR}}(x_s, \omega_o) = \int_{\Omega} L_i(\omega_i) \cdot R_s(x_s, \omega_i, \omega_o, \mathbf{n}_s) \cdot V_s(x_s, \omega_i) \cdot (\omega_i \cdot \mathbf{n}_s) d\omega_i, \quad (3)$$

where x_s denotes the surface intersection point with the incoming light ray $L_i(\omega_i)$, and \mathbf{n}_s denotes surface normal $\bar{\mathbf{n}}_{\text{surf}}$. The incoming light ray $L_i(\omega_i)$ is determined by the environment illumination. The term $R_s(x_s, \omega_i, \omega_o, \mathbf{n}_s)$ represents the BRDF, encapsulating surface material information. The visibility term $V_s(x_s, \omega_i)$ models occlusions and incidence angles between the lighting source $\{L_i\}_i$ and surface point x_s . These geometric and material components contribute to the physically-based rendering process.

Geometry Components. We explicitly model the environment illumination with an array of $16 \times 32 \times 3$ learnable light probes $\{L_i\}_i$, uniformly distributed across a surrounding sphere to simulate the environment lighting. With the refined dynamic mesh G_{refined} , we compute visibility via ray-casting on explicit geometry using Open3D (Zhou, Park, and Koltun 2018), which is typically more efficient than ray-tracing in implicit volumetric fields (e.g., (Chen and Liu 2022; Xu et al. 2024)) that require iterative field queries. And the interaction between lighting and surface normal is assessed by the incidence angle between ω_i and \mathbf{n}_s . In our case, \mathbf{n}_s is extracted from N_{surf} .

Material Component. We adopt the Micro-facet model (Walter et al. 2007) to simulate the BRDF, with the Cook-Torrance kernel (Cook and Torrance 1982) to partition surface reflection into diffuse and specular components:

$$R_s = k_{\text{diffuse}} \cdot \alpha_s + k_{\text{specular}} \cdot \gamma_s, \quad (4)$$

where k_{diffuse} and k_{specular} denote the fractions of incident light energy allocated to the diffuse and specular components, both set to 1. The diffusive component α_s is queried from the Albedo Network \mathcal{M}_{alb} , while the specular component γ_s is derived from the Roughness Network \mathcal{M}_{rgh} . Both \mathcal{M}_{alb} and \mathcal{M}_{rgh} are decoded from the hash features from $\mathcal{M}_{\text{color}}$. By efficiently combining these components, we compute the final physically-based rendering result $I_{\text{PBR}}(x_s, \omega_o)$. We then minimize the difference between I_{PBR} and I_{GT} to optimize light probes and material networks.

Joint Material-Geometry Refinement Through Generative Inverse Shading

Once the geometry and material networks produce reliable surface normals and reflectance estimates, we introduce a de-shading module $\mathcal{M}_{\text{deshade}} : (\bar{\mathbf{n}}_{\text{surf}}, \alpha_s) \rightarrow \hat{\alpha}_s$, designed using a three-level U-Net structure, to further improve the predicted albedo by removing residual shading. These shading artifacts typically result from inaccuracies in geometry, which lead to lighting being baked into the albedo.

This module leverages inverse shading priors learned from data introduced in the next section to disentangle intrinsic reflectance from shading. Rather than treating albedo and geometry separately, $\mathcal{M}_{\text{deshade}}$ jointly analyzes

the shading-influenced albedo α_s and surface normal N_{surf} to estimate a clean, unshaded albedo $\hat{\alpha}_s$. The refined albedo serves both as a physically plausible output and as a learning target for the albedo network \mathcal{M}_{alb} . Next, we compute a refined PBR image \hat{I}_{PBR} from the unshaded albedo $\hat{\alpha}_s$, replacing the original I_{PBR} . This refined image is then compared with the ground truth I_{GT} to update other components: surface normals via $\mathcal{M}_{\text{offset}}$, roughness via \mathcal{M}_{rgh} , and lighting $\{L_i\}_i$. As surface geometry improves, de-shading becomes more reliable, leading to more accurate albedo and further enhancing geometry—a positive feedback loop that refines both shape and appearance jointly.

Experiments

Datasets and Evaluation Metrics

We train and evaluate on eight datasets spanning synthetic and real humans in both indoor and outdoor settings. For synthetic humans, *SyntheticHuman* (Peng et al. 2024b) provides 7 dynamic human models with ground truth meshes. We select four subjects and four views to assess mesh reconstruction accuracy, provided in the extended version. *SyntheticHuman++* (Xu et al. 2024) offers 6 dynamic models with relighting data. We use four individuals and four views for training and comparison with relighting methods. For real-world scenarios, *People Snapshot* (Alldieck et al. 2018) contains videos of humans moving in circles. We use two outdoor sequences for qualitative evaluation with Relighting4D. *MobileStage* (Xu et al. 2024) includes multi-view images of a person’s motion, which we use for training and comparison with RA. To further analyze our approach via ablation studies, we randomly choose one sample from *ZJU-MoCap* (Peng et al. 2021), one from *THuman2.0* (Yu et al. 2021), two from *SyntheticHuman++*, and one from *People Snapshot*. For model training, we create a *Mesh Down-sampling Dataset* with 526 meshes and 35,800 color-normal image pairs, generated with the human vision foundation model Sapiens (Khirodkar et al. 2024), to train our normal enhancement diffusion model $\mathcal{M}_{\text{enhance}}$. *De-shading Dataset* is constructed using clothing patterns from (Medeiros et al. 2017) and synthetic rendering under varied lighting, providing paired data to train our de-shading module $\mathcal{M}_{\text{deshade}}$. Further details are presented in the extended version.

For surface detail evaluation, we use **Normal Degree** for normal image cosine similarity. For image quality, we report **LPIPS*** (Zhang et al. 2018) ($\text{LPIPS} \times 10^3$) and **PSNR**.

Relightable	Method	Memory (GB) ↓	Speed (FPS) ↑
No	NDF	5.5	2.48
	AniSDF	11.4	0.93
	3DGS-Avatar	4.0	50+
Yes	Relighting4D	7.2	0.33
	RA	10.89	0.52
	Ours (DIS)	3.8	12.16

Table 1: Rendering Memory and Speed Comparison. **Best** and **second-best** results for each metric are highlighted.

	Normal	Diffuse Albedo		Visibility		Relighting		Speed	
	Degree ↓	PSNR ↑	LPIPS* ↓	PSNR ↑	LPIPS* ↓	PSNR ↑	LPIPS* ↓	Training ↓	Rendering ↑
NeRFactor (1 frame)	-	22.23	226.25	11.37	387	21.04	313	50+ hours	0.3- FPS
Relighting4D	27.41	36.16	36.29	24.79	26.39	32.27	30.40	40+ hours	0.3- FPS
RA	17.11	34.08	35.36	33.97	21.07	34.59	26.21	50+ hours	0.5- FPS
Ours	15.89	36.39	34.71	35.11	19.70	35.32	25.10	3 hours	12+ FPS

Table 2: Comparison of Relighting Metrics. **Best** and **second-best** results for each metric are highlighted.

Comparisons on Rendering Efficiency

Images are rendered at 512×512 on a server with 2 Xeon Silver 4114 CPUs and 1 V100 GPU. Table 1 reports average memory and speed over 100 runs, showing that DIS requires the least memory among all methods by deforming a few vertices and using a network for surface details. In contrast, implicit approaches such as NDF (Zhang and Chen 2022) and AniSDF (Peng et al. 2024b) use coordinate-based networks to cover the entire 3D space, while explicit methods like 3DGS-Avatar (Qian et al. 2024) store a dense array of 3D Gaussians, both leading to high memory costs. DIS also supports PBR and achieves fast rendering via differentiable rasterization, while methods, such as Relighting4D and RA, are constrained by the inefficiency of coordinate-based networks. Although our rendering speed is slower than that of 3DGS methods, our accurate and detailed mesh surfaces offer substantial advantages for surface material learning, particularly when compared to sparse explicit methods that struggle to deliver similar quality.

Evaluation on Relighting

We compare our DIS with leading relighting methods, RA and Relighting4D, and NeRFactor, focusing on static objects trained from a single frame provided in (Xu et al. 2024). Our evaluation results are summarized in Table 2 and detailed in the extended version. To resolve scale ambiguity in inverse rendering, we align predicted images to ground truth using channel-wise scale factors (Zhang et al. 2021). All metrics are reported for full images and rendering speeds at 512×512 resolution on 1 NVIDIA V100 GPU, while RA is evaluated only on foreground regions. By integrating a refined mesh with detailed surface variations from normal enhancement priors, our method achieves superior normal accuracy and visibility compared to implicit approaches, resulting in the best relighting performance. Our de-shading module further yields the highest albedo metrics. We also report the shortest training time and fastest rendering speed, exceeding 12 FPS, demonstrating significant efficiency gains. As shown in Fig. 3, visual comparisons under four novel lighting conditions indicate that Relighting4D often suffers from inaccurate normals and visibility due to unreliable geometry, while RA captures fine details but is prone to instability from implicit fields. In contrast, our method consistently delivers the most stable and realistic results.

To evaluate our method’s effectiveness in disentangling real human material attributes, we qualitatively compare with Relighting4D using *People Snapshot*, following their training setup. While RA claims to work well with monocu-

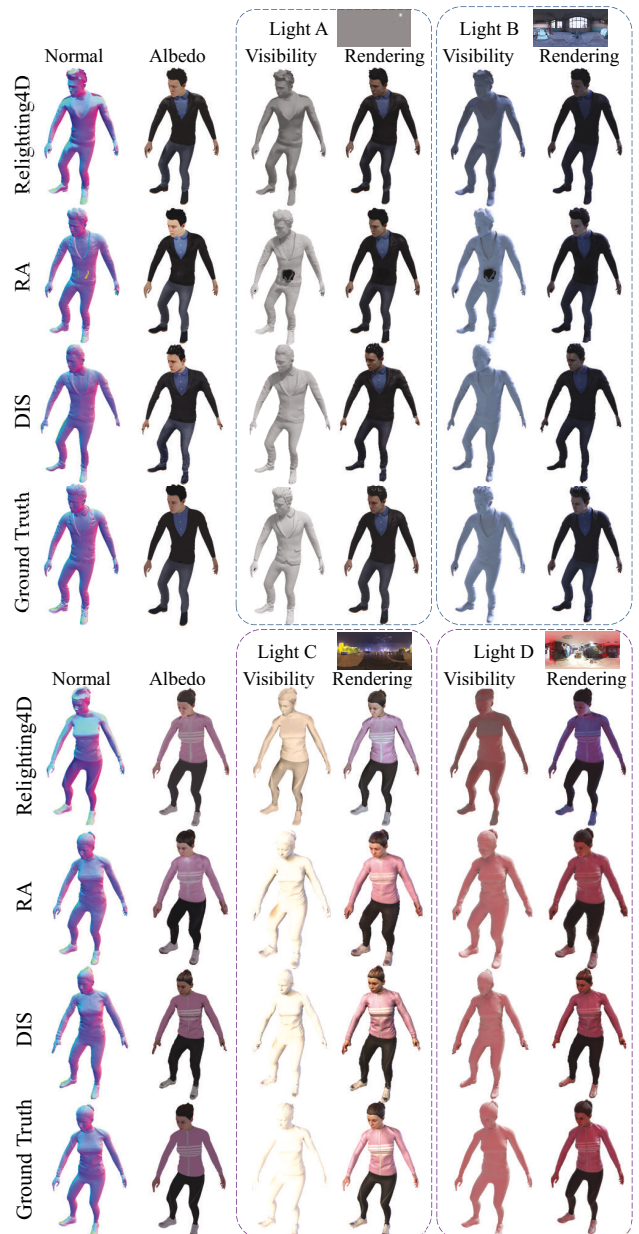


Figure 3: Qualitative comparison on *SyntheticHuman++* across DIS, Relighting4D, and RA.

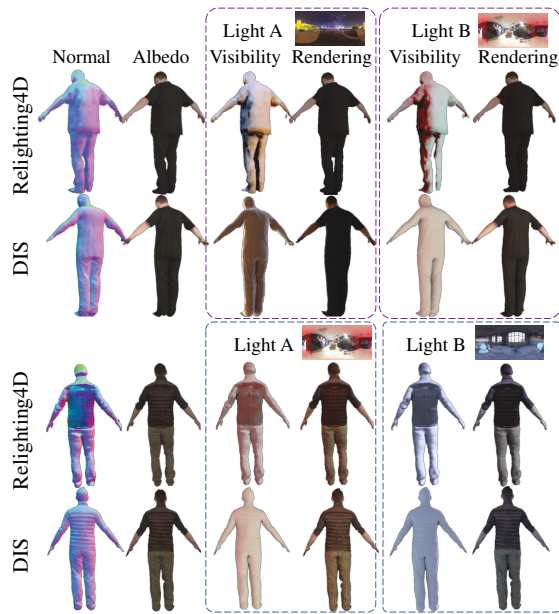


Figure 4: Qualitative comparison on *People Snapshot* (real-captured outdoor humans) between DIS and Relighting4D.

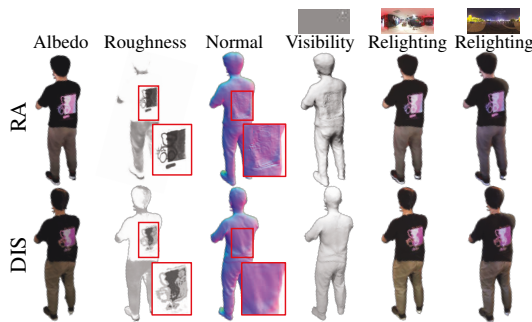


Figure 5: Qualitative comparison on *MobileStage* (real-captured indoor human) between DIS and RA.

lar input, our experiments suggest otherwise, revealing significant challenges in training relightable human avatars in real-world scenarios. In contrast, our method consistently produces good results, aided by normal enhancement priors. As shown in Fig. 4, our mesh-based representation produces superior surface normals and enables more effective material disentanglement than Relighting4D, resulting in more realistic visibility and relighting across diverse lighting conditions. And our albedo is more accurate with less shading effects, benefiting from our inverse shading priors.

We also conduct qualitative comparisons with RA on *MobileStage*, using their provided model. As shown in Fig. 5, our method, benefiting from normal enhancement priors, produces smoother surfaces and more accurate normal maps, resulting in better normal and material disentanglement. For example, the texture on the back of the clothing in our results correctly reflects material differences rather than geometric structure. Our relighting outputs are also more

realistic: under the third lighting condition, DIS captures a warm yellow light effect, whereas RA’s output appears blue. Since RA does not provide specific training settings, we randomly select 500 frames from eight views for training.

Ablation Studies

We show three ablation studies on two main components: the normal conversion module and the de-shading module.

Normal Conversion Module v.s. Mesh Split

To evaluate the mesh surface expression capability of our normal conversion module $O2N$, we compare our DIS (*w/o* $O2N$) & DIS (*w* $O2N$) with an SMPL-split-based approach *Mesh Split* that increases vertex count from 6,890 to 27,554 under identical training. As shown in Fig. 6 of the extended version, DIS (*w* $O2N$) captures finer surface details, such as wrinkles in pants, overcoming topology limitations and achieving ultra-resolution mesh quality. Table 3 of the extended version shows that *Mesh Split* requires three times longer training, while DIS (*w* $O2N$) achieves a significantly lower normal degree, demonstrating effective surface detail for PBR with efficient training.

PBR w/wo Normal Conversion Module

To highlight the importance of the normal conversion module $O2N$ in PBR, we compare visibility results with and without $O2N$ in Fig. 7 of the extended version. The red rectangles show that incorporating $O2N$ allows the visibility map to capture fine details, such as wrinkles on the tie and clothing pockets. This also enhances lighting interactions, resulting in sharper edges and rendering results that more closely match GT.

Albedo & Normal w/wo De-shading Module

To demonstrate the effectiveness of our de-shading module, Fig. 8 of the extended version visualizes the de-shaded albedo for both a synthetic and a real human. As highlighted in red boxes, our de-shading module effectively removes clothing reflections in synthetic humans and facial glare in real humans, demonstrating strong albedo restoration. This capability explains our SOTA albedo accuracy. Additionally, Table 4 of the extended version shows a decrease in normal degree after applying inverse shading priors, further validating our approach.

Conclusion

We introduce DIS for dynamic and relightable avatar modeling, distilling normal enhancement and inverse shading priors onto 3D surfaces through a normal conversion module. Using differential rasterization between 2D images and 3D surfaces, DIS effectively disentangles materials from sparse-view video input, achieving SOTA relighting results and efficient rendering. Joint optimization of generative priors for geometry, material, and surface normals further enhances surface quality and material accuracy.

Acknowledgments

This research was supported by the Theme-based Research Scheme, Research Grants Council of Hong Kong (T45-205/21-N), and the Guangdong and Hong Kong Universities “1+1+1” Joint Research Collaboration Scheme (2025A0505000003).

References

- Alldieck, T.; Magnor, M.; Xu, W.; Theobalt, C.; and Pons-Moll, G. 2018. Video based reconstruction of 3d people models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8387–8397.
- Chen, Z.; and Liu, Z. 2022. Relighting4D: Neural Relightable Human from Videos. *European Conference on Computer Vision*.
- Chen, Z.; Moon, G.; Guo, K.; Cao, C.; Pidhorskyi, S.; Simon, T.; Joshi, R.; Dong, Y.; Xu, Y.; Pires, B.; Wen, H.; Evans, L.; Peng, B.; Buffalini, J.; Trimble, A.; McPhail, K.; Schoeller, M.; Yu, S.-I.; Romero, J.; Zollhöfer, M.; Sheikh, Y.; Liu, Z.; and Saito, S. 2024. URHand: Universal Relightable Hands. [arXiv:2401.05334](https://arxiv.org/abs/2401.05334).
- Cook, R. L.; and Torrance, K. E. 1982. A reflectance model for computer graphics. *ACM Transactions on Graphics (ToG)*, 1(1): 7–24.
- Debevec, P.; Hawkins, T.; Tchou, C.; Duiker, H.-P.; Sarokin, W.; and Sagar, M. 2000. Acquiring the reflectance field of a human face. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, 145–156.
- Dou, Y.; Zheng, Z.; Jin, Q.; Shi, R.; Li, Y.; and Ni, B. 2024. Differentiable Micro-Mesh Construction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4294–4303.
- Gafni, G.; Thies, J.; Zollhofer, M.; and Nießner, M. 2021. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8649–8658.
- Guo, K.; Lincoln, P.; Davidson, P.; Busch, J.; Yu, X.; Whalen, M.; Harvey, G.; Orts-Escolano, S.; Pandey, R.; Dourgarian, J.; et al. 2019. The relightables: Volumetric performance capture of humans with realistic relighting. *ACM Transactions on Graphics*, 38(6): 1–19.
- Habermann, M.; Xu, W.; Zollhoefer, M.; Pons-Moll, G.; and Theobalt, C. 2019. Livecap: Real-time human performance capture from monocular video. *ACM Transactions On Graphics (TOG)*, 38(2): 1–17.
- He, Z.; Wang, T.; Huang, X.; Pan, X.; and Liu, Z. 2024. Neural LightRig: Unlocking Accurate Object Normal and Material Estimation with Multi-Light Diffusion. [arXiv:2412.09593](https://arxiv.org/abs/2412.09593).
- Hu, L.; Gao, X.; Zhang, P.; Sun, K.; Zhang, B.; and Bo, L. 2023. Animate Anyone: Consistent and Controllable Image-to-Video Synthesis for Character Animation. [arXiv:2311.17117](https://arxiv.org/abs/2311.17117).
- Huang, X.; Shao, R.; Zhang, Q.; Zhang, H.; Feng, Y.; Liu, Y.; and Wang, Q. 2024a. Humannorm: Learning normal diffusion model for high-quality and realistic 3d human generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4568–4577.
- Huang, Y.; Yi, H.; Xiu, Y.; Liao, T.; Tang, J.; Cai, D.; and Thies, J. 2024b. TeCH: Text-guided Reconstruction of Life-like Clothed Humans. In *International Conference on 3D Vision (3DV)*.
- Iqbal, U.; Caliskan, A.; Nagano, K.; Khamis, S.; Molchanov, P.; and Kautz, J. 2023. Rana: Relightable articulated neural avatars. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 23142–23153.
- Jin, H.; Liu, I.; Xu, P.; Zhang, X.; Han, S.; Bi, S.; Zhou, X.; Xu, Z.; and Su, H. 2023. TensorIR: Tensorial Inverse Rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4): 1–14.
- Khrodar, R.; Bagautdinov, T.; Martinez, J.; Zhaoen, S.; James, A.; Selednik, P.; Anderson, S.; and Saito, S. 2024. Sapiens: Foundation for Human Vision Models. [arXiv preprint arXiv:2408.12569](https://arxiv.org/abs/2408.12569).
- Kim, H.; Jang, M.; Yoon, W.; Lee, J.; Na, D.; and Woo, S. 2024. SwitchLight: Co-design of Physics-driven Architecture and Pre-training Framework for Human Portrait Relighting. [arXiv:2402.18848](https://arxiv.org/abs/2402.18848).
- Li, J.; Cao, C.; Schwartz, G.; Khrodar, R.; Richardt, C.; Simon, T.; Sheikh, Y.; and Saito, S. 2024. URAvatar: Universal Relightable Gaussian Codec Avatars. [arXiv:2410.24223](https://arxiv.org/abs/2410.24223).
- Li, P.; Zheng, W.; Liu, Y.; Yu, T.; Li, Y.; Qi, X.; Chi, X.; Xia, S.; Cao, Y.-P.; Xue, W.; Luo, W.; and Guo, Y. 2025. PSHuman: Photorealistic Single-image 3D Human Reconstruction using Cross-Scale Multiview Diffusion and Explicit Remeshing. [arXiv:2409.10141](https://arxiv.org/abs/2409.10141).
- Long, X.; Guo, Y.-C.; Lin, C.; Liu, Y.; Dou, Z.; Liu, L.; Ma, Y.; Zhang, S.-H.; Habermann, M.; Theobalt, C.; et al. 2023. Wonder3D: Single Image to 3D using Cross-Domain Diffusion. [arXiv preprint arXiv:2310.15008](https://arxiv.org/abs/2310.15008).
- Loper, M.; Mahmood, N.; Romero, J.; Pons-Moll, G.; and Black, M. J. 2015. SMPL: A skinned multi-person linear model. *ACM transactions on graphics*, 34(6): 1–16.
- Maggiordomo, A.; Moreton, H.; and Tarini, M. 2023. Micro-mesh construction. *ACM Transactions on Graphics (TOG)*, 42(4): 1–18.
- Medeiros, A. J.; Stearns, L.; Findlater, L.; Chen, C.; and Froehlich, J. E. 2017. Recognizing Clothing Colors and Visual Textures Using a Finger-Mounted Camera: An Initial Investigation. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility - ASSETS '17*.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2020. NeRF: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, 405–421.

- Müller, T.; Evans, A.; Schied, C.; and Keller, A. 2022. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4): 1–15.
- Pavlakos, G.; Choutas, V.; Ghorbani, N.; Bolkart, T.; Osman, A. A. A.; Tzionas, D.; and Black, M. J. 2019. Expressive Body Capture: 3D Hands, Face, and Body from a Single Image. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Peng, B.; Wang, J.; Zhang, Y.; Li, W.; Yang, M.-C.; and Jia, J. 2024a. ControlNeXt: Powerful and Efficient Control for Image and Video Generation. *arXiv preprint arXiv:2408.06070*.
- Peng, S.; Xu, Z.; Dong, J.; Wang, Q.; Zhang, S.; Shuai, Q.; Bao, H.; and Zhou, X. 2024b. Animatable implicit neural representations for creating realistic avatars from videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Peng, S.; Zhang, Y.; Xu, Y.; Wang, Q.; Shuai, Q.; Bao, H.; and Zhou, X. 2021. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *IEEE Conference on Computer Vision and Pattern Recognition*, 9054–9063.
- Poole, B.; Jain, A.; Barron, J.; and Mildenhall, B. 2022. DreamFusion: Text-to-3D using 2D Diffusion. *International Conference on Learning Representations*.
- Pumarola, A.; Corona, E.; Pons-Moll, G.; and Moreno-Noguer, F. 2021. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10318–10327.
- Qian, Z.; Wang, S.; Mihajlovic, M.; Geiger, A.; and Tang, S. 2024. 3DGS-Avatar: Animatable Avatars via Deformable 3D Gaussian Splatting.
- Ravi, N.; Reizenstein, J.; Novotny, D.; Gordon, T.; Lo, W.-Y.; Johnson, J.; and Gkioxari, G. 2020. Accelerating 3D Deep Learning with PyTorch3D. *arXiv:2007.08501*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. *CVPR*.
- Shu, Z.; Yumer, E.; Hadap, S.; Sunkavalli, K.; Shechtman, E.; and Samaras, D. 2017. Neural face editing with intrinsic image disentangling. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5541–5550.
- Srinivasan, P. P.; Deng, B.; Zhang, X.; Tancik, M.; Mildenhall, B.; and Barron, J. T. 2021. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7495–7504.
- Sun, W.; Che, Y.; Huang, H.; and Guo, Y. 2023. Neural reconstruction of relightable human model from monocular video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 397–407.
- Voleti, V. S.; Yao, C.-H.; Boss, M.; Letts, A.; Pankratz, D.; Tochilkin, D.; Laforte, C.; Rombach, R.; and Jampani, V. 2024. SV3D: Novel Multi-view Synthesis and 3D Generation from a Single Image using Latent Video Diffusion. *European Conference on Computer Vision*.
- Walter, B.; Marschner, S. R.; Li, H.; and Torrance, K. E. 2007. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*, 195–206.
- Wang, P.; Liu, L.; Liu, Y.; Theobalt, C.; Komura, T.; and Wang, W. 2021. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*.
- Wang, S.; Schwarz, K.; Geiger, A.; and Tang, S. 2022. Arah: Animatable volume rendering of articulated human sdfs. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXII*, 1–19. Springer.
- Wu, J.; Zhang, R.; Chen, J.; and Zhang, H. 2025. Fast and Physically-based Neural Explicit Surface for Relightable Human Avatars. *arXiv:2503.18408*.
- Xiu, Y.; Ye, Y.; Liu, Z.; Tzionas, D.; and Black, M. J. 2024. PuzzleAvatar: Assembling 3D Avatars from Personal Albums. *arXiv:2405.14869*.
- Xu, Z.; Peng, S.; Geng, C.; Mou, L.; Yan, Z.; Sun, J.; Bao, H.; and Zhou, X. 2024. Relightable and Animatable Neural Avatar from Sparse-View Video. In *CVPR*.
- Yu, T.; Zheng, Z.; Guo, K.; Liu, P.; Dai, Q.; and Liu, Y. 2021. Function4D: Real-time Human Volumetric Capture from Very Sparse Consumer RGBD Sensors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR2021)*.
- Zhang, K.; Luan, F.; Wang, Q.; Bala, K.; and Snavely, N. 2021. Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5453–5462.
- Zhang, R.; and Chen, J. 2022. NDF: Neural Deformable Fields for Dynamic Human Modelling. In *European Conference on Computer Vision*, 1–14. Springer.
- Zhang, R.; and Chen, J. 2024. Mesh-Centric Gaussian Splatting for Human Avatar Modelling with Real-time Dynamic Mesh Reconstruction. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 6823–6832.
- Zhang, R.; Chen, J.; and Wang, Q. 2023. Explicifying Neural Implicit Fields for Efficient Dynamic Human Avatar Modeling via a Neural Explicit Surface. In *Proceedings of the 31st ACM International Conference on Multimedia*, 1955–1963.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zhou, Q.-Y.; Park, J.; and Koltun, V. 2018. Open3D: A Modern Library for 3D Data Processing. *arXiv:1801.09847*.
- Zhu, S.; Chen, J. L.; Dai, Z.; Xu, Y.; Cao, X.; Yao, Y.; Zhu, H.; and Zhu, S. 2024. Champ: Controllable and Consistent Human Image Animation with 3D Parametric Guidance. In *European Conference on Computer Vision (ECCV)*.