

Learning Implicit Body Representations from Double Diffusion Based Neural Radiance Fields

Guangming Yao¹, Hongzhi Wu¹, Yi Yuan^{2*}, Lincheng Li^{2†}, Kun Zhou¹ and Xin Yu³

¹Zhejiang University

²NetEase Fuxi AI Lab

³University of Technology Sydney

Abstract

In this paper, we present a novel double diffusion based neural radiance field, dubbed DD-NeRF, to reconstruct human body geometry and render the human body appearance in novel views from a sparse set of images. We first propose a double diffusion mechanism to achieve expressive representations of input images by fully exploiting human body priors and image appearance details at two levels. At the coarse level, we first model the coarse human body poses and shapes via an unclothed 3D deformable vertex model as guidance. At the fine level, we present a multi-view sampling network to capture subtle geometric deformations and image detailed appearances, such as clothing and hair, from multiple input views. Considering the sparsity of the two level features, we diffuse them into feature volumes in the canonical space to construct neural radiance fields. Then, we present a signed distance function (SDF) regression network to construct body surfaces from the diffused features. Thanks to our double diffused representations, our method can even synthesize novel views of unseen subjects. Experiments on various datasets demonstrate that our approach outperforms the state-of-the-art in both geometric reconstruction and novel view synthesis.

1 Introduction

Simultaneous reconstruction of body geometry and appearances from a sparse set of views is highly challenging yet important in a wide variety of applications, including special effects, game production and virtual reality. Substantial research efforts have been made in the past, particularly with the help of powerful deep learning techniques. When high-quality 3D body scans are available for network training, deep networks can reconstruct human body geometry and texture easily. However, such data are often expensive to acquire and not publicly accessible.

*Corresponding Author

†Corresponding Author




Figure 1: Our method produces detailed body geometry and synthesizes high-fidelity novel views, from a sparse set of person images.

When 3D supervision is not available, self-supervised neural radiance fields (NeRF) have been proposed to render novel views of a specific subject. However, this per-subject optimization process of NeRF is often time-consuming. Recent works [Yu *et al.*, 2021; Chen *et al.*, 2021; Wang *et al.*, 2021b] propose to first extract features of scene images and then utilize a scene-shared NeRF to synthesize novel views from the scene features, thus circumventing the inefficient per-subject optimization. However, these methods would suffer artifacts when they are directly applied to human body reconstruction or rendering due to the non-rigid deformations and large variations of human bodies. As suggested in [Peng *et al.*, 2021], the introduction of human body priors will significantly alleviate geometric distortions. Unfortunately, the work [Peng *et al.*, 2021] is a per-subject optimization method, thus restricting its applications in rendering various subjects.

To overcome the aforementioned shortcomings, we present a novel double diffusion based neural radiance field, namely DD-NeRF. As illustrated in Fig. 1, DD-NeRF is developed to reconstruct human body geometry and render human body appearances in novel views from a sparse set of human body images. To be specific, we propose a novel double diffusion mechanism that can fully exploit human body priors as well as image appearance details at two levels to achieve representative features of input images.

At the coarse level, we first model the coarse human body poses and shapes via an unclothed deformable human mesh as explicit guidance. We embed the deformed vertices as a coarse-level feature, which indicates the coarse-scale human body pose and shape. At the fine level, we develop a multi-view sampling network to capture image appearance details and subtle geometric deformations, such as clothing and hair, which are not modeled at the coarse level. Our multi-view sampling network firstly samples 3D points on the deformable human model as anchors and then projects them to each input view. Then, it aggregates image features at the project positions of different views by measuring their means and variances as the fine level features. In this fashion, the acquired coarse level and fine level feature representations are both sparse in the canonical space. Afterwards, we diffuse them into 3D feature volumes via sparse convolutions to facilitate the decoding process.

Subsequently, we design an implicit field regression module to reconstruct high-quality human bodies from the two level diffused features. Here, we employ a signed distance function (SDF) to represent body geometry instead of volume-density based radiance fields since SDF is accurate in modeling surface geometry while volume density is more suitable to render texture [Wang *et al.*, 2021a]. Thus, we adopt multi-layer perceptrons (MLPs) to regress signed distance from the tri-linearly interpolated coarse and fine level feature volumes. In order to avoid the per-subject optimization as in [Peng *et al.*, 2021], our radiance regression network is conditioned on image features and pixels of input views in reasoning. Additionally, we employ a transformer to fuse the multi-view image features and raw pixels into a joint representation for radiance regression. With regressed signed distance and radiance, the reconstructed meshes and novel views can be obtained from differential SDF rendering [Wang *et al.*, 2021a]. Extensive experiments demonstrate that once DD-NeRF is trained, it can reconstruct human bodies of both seen and unseen subjects and achieves superior performance compared to the state-of-the-art.

In summary, the contributions of this work are three-fold:

- We present a novel double diffusion based neural radiance field, dubbed DD-NeRF, which can reconstruct human body geometry from a sparse set of images in a feed-forward fashion and is subject-agnostic.
- We propose a double diffusion mechanism that fully exploits human body priors and captures image appearance details in both coarse and fine levels, enabling accurate and robust human body surface reconstruction.
- We develop a signed distance function (SDF) regression network to reconstruct human body surface in a differentiable manner, allowing the body surface to be optimized with only 2d supervision.

2 Related Work

2.1 Body Shape Reconstruction

Recently, regressing a “freeform” 3D body shape with implicit representations has achieved promising progress. Pixel-aligned implicit function (PIFu) regresses a signed distance

function for any given 3D location [Saito *et al.*, 2019; Saito *et al.*, 2020]. It can infer both 3D body surface and texture from a single image. Several works take multi-view images as input to pursue better reconstruction performance. For instance, volumetric occupancy fields have been proposed to learn dynamic clothed bodies from sparse viewpoints [Gilbert *et al.*, 2018]. The work [Shao *et al.*, 2021] combines a surface field and a radiance field for body representations, and it is learned under the supervision of ground-truth meshes. Note that all these methods require ground-truth 3D scanned models as supervision. In contrast, our method does not rely on 3D scanned models and only leverages a sparse set of multi-view images to reconstruct high-quality 3D body geometry in a self-supervised manner.

2.2 Neural Representations

Neural implicit functions have emerged as an effective representation to learn 3D scenes from 2D images. NeRF [Mildenhall *et al.*, 2020] learns densities and colors of a scene with volumetric rendering, and achieves impressive results on novel view synthesis. Following NeRF, some methods [Peng *et al.*, 2021; Kwon *et al.*, 2021] capture a human body geometry from videos and reconstruct its surface by performing marching cube [Lorensen and Cline, 1987] on learned volumetric densities. Furthermore, several works [Niemeyer *et al.*, 2020; Yariv *et al.*, 2020; Kellnhofer *et al.*, 2021; Wang *et al.*, 2021a] represent a scene with an SDF and thus extract surfaces by level set. The work [Peng *et al.*, 2021] employs a skinned multi-person linear model (SMPL) [Loper *et al.*, 2015] as a human body prior to preserve geometric details of human bodies, but it needs to optimize the network according to each subject. Recent works [Chen *et al.*, 2021; Wang *et al.*, 2021b] propose conditional NeRF to bypass the tedious per-subject optimization. They are trained on multiple subjects and can perform novel view synthesis for unseen subjects which are not available during training. Our method also learns across different subjects in order to reconstruct unseen subjects in a feed-forward fashion. Moreover, we introduce a human body model as geometric guidance, thus significantly facilitating high-fidelity body geometry reconstruction and improving optimization efficiency.

3 Proposed Method

Given a sparse set of multi-view posed body images, our method computes an implicit radiance field that represents the geometry and appearance of the subject. We denote the input images as $I = \{I_1, I_2, \dots, I_{N_c}\}$, where N_c is the number of pre-calibrated cameras. The calibration results are stored as $\Phi = \{\Phi_1, \Phi_2, \dots, \Phi_{N_c}\}$. In general, our network can be viewed as a conditional radiance field function as follows:

$$s, c = F(x, d; I, \Phi), \quad (1)$$

where x represents a 3D location, d is a viewing direction, s denotes a signed distance value at x and c is an RGB color as its appearance. The output radiance field can be used to synthesize a novel view via differentiable ray marching, or extract 3D surfaces with marching cube.

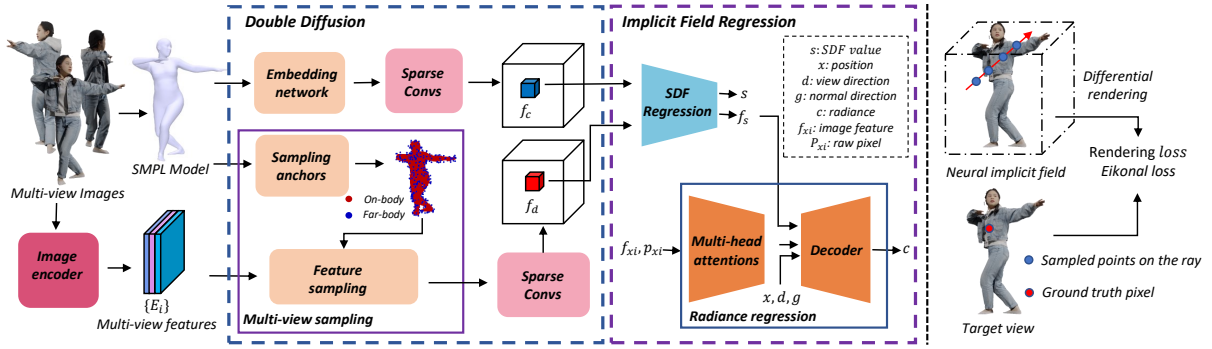


Figure 2: An overview of *DD-NeRF*. Given multi-view input images and a corresponding SMPL model, the double diffusion module constructs two feature volumes to represent a coarse body prior and local features that depict local geometric deformations. Next, an implicit field regression module predicts a signed distance function (SDF) and scene radiance for each ray. All the modules are learned jointly by minimizing the differences between the target and rendered pixels.

An overview of our model is illustrated in Fig. 2. There are three components: an image encoder module, a double diffusion module, and an implicit field regression module. The image encoder module is composed of a stacked hourglass neural network [Newell *et al.*, 2016] for image feature extraction. Next, the double diffusion module builds two volumes, a coarse body feature volume and a fine detail feature volume. While the former is designed to encode the coarse body prior, the latter is developed to depict the detailed geometric deformations for high-fidelity body representation. The implicit field regression module consists of two sub-modules that regress signed distance values and colors, respectively. Specifically, an SDF regression sub-module generates signed distance from diffused features of both volumes, and a radiance regression sub-module takes as input surface normal, multi-view raw pixels, image features and the view direction, and produces an RGB color as output. Once we obtain the regressed signed distance and colors, a novel view image of the subject can be synthesized by differential SDF rendering.

3.1 Double Diffusion

The motivation of double diffuse is to attain expressive representations by exploiting both coarse human body priors and detailed multi-view image features. We build a coarse body volume and a fine-level detail feature volume in parallel. We use the resolution of 256^3 for the coarse body volume, and 512^3 for the detail feature volume in order to capture high-frequency geometric details. Inspired by [Peng *et al.*, 2021], we adopt an SMPL model as the body prior and the coarse body volume encodes the body shapes and poses. Given the multi-view images, the corresponding SMPL parameters are estimated by EasyMocap [Dong and Fang, 2021]. We obtain the vertices $V = \{v_1, v_2, \dots, v_{6,890}, v_i \in R^3\}$ from the template model. The vertices are fed into the embedding network to compute a structured latent code $Z \in R^{6,890 \times 16}$, which serves as our coarse body prior for implicit field reconstruction.

The detail feature volume is designed to capture fine-grained geometric characteristics of a subject by integrating multi-view image features. Inspired by [Yu *et al.*, 2018], we sample the on-body and far-body anchors to effectively

capture the body-related features. The on-body anchors are relatively dense vertices computed with subdivisions on the SMPL model. The far-body anchors are randomly sampled on the exterior surfaces of spheres centered at each vertex of the SMPL model. We empirically set the radius as 0.05m to account for objects that are relatively far away from the SMPL surface (e.g., loose clothing, long hair). We uniformly sample 12,446 far-body anchors, and obtain 27,554 on-body anchors from 6,890 SMPL vertices with subdivision. Next, we project the anchors to different views and sample pixel-aligned multi-view features, expressed as:

$$\begin{aligned} \hat{a}_i^j &= K_j \cdot [R_j^T \cdot R \cdot (a_i - T) + T_j], \\ f_i^j &= E_j\{\hat{a}_i^j\}, \end{aligned} \quad (2)$$

where a_i is the 3D position of the i -th anchor, R and T are the estimated rotation and translation of the SMPL model. K_j and $[R_j, T_j]$ are the intrinsic and extrinsic matrices of j -th camera respectively. $E_j\{\cdot\}$ denotes the image feature extracted from the j -th view, $f_i^j \in R^c$ denotes the sampled features of the i -th anchor from the j -th view. Given the N_c input views, we attain features $f_i = \{f_1, f_2, \dots, f_{N_c}\}$ for each anchor and then concatenate their mean and variance to achieve a feature vector $\hat{f}_i \in R^{2c}$ as the final representation of the anchor. Then, those feature vectors are utilized to capture subtle geometric deformations and image detailed appearances.

Since the sparse coarse-and-fine features are not suitable for dense surface reconstruction, we employ SparseConvNet [Graham and Engelcke, 2018] to diffuse them into the canonical space. Our double-diffusion module incorporates the coarse body prior with elaborated geometric information while preserving the local appearance and geometric details depicted in input images.

3.2 Implicit Field Regression

This step consists of two sub-modules: SDF regression F_S and radiance regression F_R . To obtain a high-quality output surface, we use signed distance instead of volumetric density to represent the body geometry. For a 3D point x that a ray traverses in the 3D bounding box Ω^3 , we trilinearly interpolate features from the coarse and fine level feature volumes

to obtain its corresponding coarse feature f_c and detail feature f_d . Then, we adopt an MLPs to regress signed distance s based on x , f_c and f_d as follows:

$$s, f_s = F_S(x, f_c, f_d), \quad (3)$$

where f_s is a surface feature corresponding to s . Similar to NeRF [Mildenhall *et al.*, 2020], x is mapped to a higher dimensional space with positional encoding. After computing s for the feature volumes, the body surface \mathcal{S} can be expressed by the zero level-set of SDF, i.e., $\mathcal{S} = \{x \in \Omega^3 | s = 0\}$.

Once the geometry is obtained, the next step is to regress the corresponding appearance. Given the geometry-related feature f_s , we take the view direction, the geometry and the image feature into consideration. The radiance regression is expressed by:

$$c = F_R(f_s, g, x, d, \{f_{xi}\}, \{p_{xi}\} | i = 1, 2, \dots, N_c), \quad (4)$$

where d is the view-direction, and $g = \nabla F_S(x, f_c, f_d)$ represents the normal of surface \mathcal{S} at position x . We project x to the i -th view, and then sample pixel-aligned image feature f_{xi} from E_i and the raw RGB pixel p_{xi} . Notably, p_{xi} and d are also mapped to a higher dimensional space with positional encoding for learning high-frequency variations.

We adopt a transformer [Vaswani and Shazeer, 2017] to effectively fuse features $\{f_{xi}\}, \{p_{xi}\} | i = 1, 2, \dots, N_c\}$ for appearance reasoning. The transformer consists of an encoder and a decoder. The encoder \mathcal{E} encodes the multi-view features with stacked multi-head attention layers to obtain a fused feature \hat{f}_x , computed as:

$$\hat{f}_x = \mathcal{E}(\{f_{xi}\}, \{p_{xi}\} | i = 1, 2, \dots, N_c). \quad (5)$$

The MLP-based decoder \mathcal{D} regresses the final radiance:

$$c = \mathcal{D}(f_s, g, x, d, \hat{f}_x). \quad (6)$$

3.3 Differential SDF Rendering

To optimize our representation, we adopt the SDF-based differential renderer [Wang *et al.*, 2021a]. Then, we can use our input images as ground-truth to supervise novel view synthesis and reconstruction. For a pixel in the target image, We accumulate the radiance along the ray emitted from the pixel to to predict the pixel color by:

$$\hat{C} = \sum_{i=1}^n T_i \alpha_i c_i, \quad (7)$$

where $T_i = \prod_{j=1}^{i-1} (1 - \alpha_j)$ is the discrete accumulated transmittance, and α_i is a discrete opacity value, defined as:

$$\alpha_i = \max\left(\frac{\Phi_s(f(s_i) - \Phi_s(f(s_{i+1})))}{\Phi_s(f(s_i))}, 0\right), \quad (8)$$

$$\Phi_s = (1 + e^{-kx})^{-1},$$

where k is a learnable scalar that increases as the training iteration progresses.

3.4 Loss Function

To learn the implicit body representations, we penalize the differences between the rendered pixel colors and their counterparts in the input image. We randomly select n input views, and randomly sample m pixels from each view to train our network. The rendering loss L_r measures the pixel-wise L1 distance between the rendered colors and the ground-truth:

$$L_r = \frac{1}{m} \sum_{i=1}^m |C_i - \hat{C}_i|, \quad (9)$$

where \hat{C}_i is the predicted pixel color and C_i is the ground-truth. Besides, the Eikonal loss L_e serves as an implicit geometric regularization [Gropp and Yariv, 2020], and enforces our SDF regression sub-module to model signed distance:

$$L_e = \frac{1}{nm} \sum_{i,j} (|\nabla(s_{ij})| - 1)^2. \quad (10)$$

Finally, the overall loss is defined as:

$$L = \lambda_r L_r + \lambda_e L_e, \quad (11)$$

where λ_r and λ_e denote the corresponding weights.

4 Experiments

4.1 Datasets

We perform experiments on the synthesized datasets Twin-dom¹, THuman2.0 [Zheng *et al.*, 2021] and the real-world dataset ZJU-Mocap [Peng *et al.*, 2021]. We use 1,200 body meshes from Twin-dom for training, 300 meshes for evaluation. For THuman, 400 meshes are used for training, and 100 meshes for evaluation. All the data in ZJU-Mocap are used for evaluation. To render a scanned body mesh, we place it in the center of a unit sphere, and orient the camera towards the sphere center with a distance of 2.4m. We move the camera around the sphere, sample a yaw angle from -30° to 60° with an interval of 10° , and sample a roll angle from 0° to 360° with an interval of 24° . For each body model, we render 135 images of resolution 1024^2 for training.

4.2 Implementation Details

Note that sampling strategies can heavily influence the final results, we adopt the hierarchical sampling method as in [Wang *et al.*, 2021a]. We use the Adam optimizer [Kingma and Ba, 2014], and set the learning rate to 1×10^{-4} . The loss weights are set to 10, 1 for λ_r and λ_e , respectively.

4.3 Qualitative Comparisons

We compare our approach with state-of-the-art methods to validate the superiority of our method. MVSNeRF [Chen *et al.*, 2021], PIFu [Saito *et al.*, 2019] and our method are all generalizable. In other words, models are trained on several subjects and tested on randomly sampled subjects that are unseen during training. NeuralBody [Peng *et al.*, 2021] is a subject-specific optimization based method. Following the training protocols of NeuralBody, we train it with 500 epochs

¹web.twindom.com

Model	Supervision	PNSR \uparrow	SSIM \uparrow	LPIPS \downarrow	PNSR \uparrow	SSIM \uparrow	LPIPS \downarrow	PNSR \uparrow	SSIM \uparrow	LPIPS \downarrow
		THuman2.0			Twindom			ZJU-Mocap		
MVSNeRF	images	18.99	0.72	0.35	16.43	0.66	0.30	19.50	0.76	0.41
MVSNeRF(Ft)	images	21.09	<u>0.80</u>	0.32	19.90	0.72	0.27	22.10	0.88	0.32
PIFu	3D meshes	18.39	0.58	0.33	16.43	0.70	0.30	19.12	0.75	0.43
NeuralBody	images	<u>21.90</u>	0.85	<u>0.20</u>	<u>23.52</u>	0.76	<u>0.26</u>	25.65	0.92	0.27
Ours	images	21.06	0.75	0.21	23.40	<u>0.79</u>	<u>0.26</u>	23.95	0.79	<u>0.25</u>
Ours(Ft)	images	23.08	0.85	0.17	27.67	0.81	0.23	<u>25.32</u>	<u>0.89</u>	0.23

Table 1: Quantitative comparisons of novel view synthesis. ‘‘Ft’’ indicates that fine-tuning is applied. Bold and underlined numbers correspond to the best and the second-best values for each metric.




Figure 3: Comparisons of geometry reconstruction results by different methods.




Figure 4: Comparisons of novel view synthesis. Our method produces the most faithful results.

Model	Chamfer \downarrow	P2S \downarrow	Chamfer \downarrow	P2S \downarrow	
		THuman2.0		Twindom	
MVSNeRF	1.597	1.146	1.528	1.126	
MVSNeRF(Ft)	0.930	0.913	0.925	0.741	
PIFu	1.510	1.524	1.170	1.630	
NeuralBody	0.915	0.931	0.815	0.725	
Ours	<u>0.810</u>	<u>0.739</u>	<u>0.803</u>	<u>0.714</u>	
Ours(Ft)	0.740	0.689	0.696	0.647	

Table 2: Quantitative comparisons of surface reconstruction.

for each testing subject. It is worth mentioning that we use the public-released weights of PIFu for comparisons. For the remaining methods, we use four views (i.e., the left, front, right and back views) as the input in testing.

We conduct comparison experiments on both the recon-

structed geometry (Fig. 3) and rendered appearances (Fig. 4). Please refer to the supplementary video for animated results with varying views. The results show that PIFu suffers over-smoothing artifacts and distortions in its reconstructed geometry and texture. Benefiting from our double diffusion mechanism, our method effectively learns both the coarse body prior and image details and thus produces the most authentic geometry and appearances. MVSNeRF and NeuralBody generate plausible novel views. However, some bumpy artifacts appear in their generated shape results as seen in Fig. 3. Since we employ SDF to represent the geometry, our results do not exhibit such artifacts. We also compare our method with NeuralBody on real-world images. Note that we only need to spend 30 minutes to finetune our method on real-world images, while NeuralBody requires 4 hours to learn from scratch for each subject on an Nvidia RTX3090 GPU.




Figure 5: Comparisons of synthesized novel views and reconstructed surfaces of real-world images.

Model	Chamfer↓	P2S↓	PSNR↑	SSIM↑
<i>w/o f_c</i>	0.843	0.91	17.03	0.70
<i>w/o DD</i>	0.926	1.08	14.54	0.49
<i>w/o transformer</i>	0.835	0.74	19.95	0.73
full model	0.803	0.71	23.40	0.79

Table 3: Ablation study on the Twindom dataset.

As illustrated in Fig. 5, our method produces better surface details and comparable novel views. For more visual results, please refer to the supplementary material.

4.4 Quantitative Comparisons

We adopt a variety of metrics to quantitatively evaluate our results from the perspectives of the image and geometry quality. Specifically, we employ peak signal-to-noise ratio (PSNR), structural similarity index (SSIM) and learned perceptual image patch similarity (LPIPS) [Zhang *et al.*, 2018] to evaluate the similarity between the ground-truth and the synthesized image. The metrics of Chamfer distance (Chamfer) and average point-to-surface Euclidean distance (P2S) are applied to geometric quality assessment. For the image quality evaluation, we randomly select 20 target views from the 135 rendered images as the ground-truth. In the geometry quality assessment, marching cube is performed for surface extraction. To further inspect the capability of generalization, we fine-tune MVSNerF and our method for 100 epochs.

The quantitative comparisons in terms of image and geometry quality are reported in Tab. 1 and 2, respectively. Our method outperforms the other competing methods on most metrics. In particular, our method achieves the highest scores at Chamfer and P2S, indicating the effectiveness of our proposed implicit body representations. In other words, our method attains the most accurate geometric models. Though DD-NeRF works in a feed-forward way during testing, we can fine-tune it, similar to the subject-specific optimization based method NeuralBody. After fine-tuning, our method achieves the highest scores in novel view synthesis. Moreover, all the metrics are also improved. This indicates that DD-NeRF can not only be extended to subject-specific novel view synthesis but also capture more expressive features.

4.5 Ablation Study

We conduct ablation studies on Twindom dataset to evaluate the contributions of each component in DD-NeRF. We first remove the coarse-level feature f_c (*w/o f_c*) to evaluate the effect of body prior. Note that the fine-level feature f_d cannot be removed since it provides all the clothes and hair information.




Figure 6: Qualitative ablation study on synthesized novel views. Our full model leads to the best result compared to other variants.

Then, we remove the double diffusion mechanism (*w/o DD*). The ablated model is trained from scratch in a subject-specific way. To inspect the effect of the transformer, we remove the transformer from the full model (*w/o transformer*), and thus multi-view features are stacked as the input of our radiance decoder. The results are reported in Tab. 3. Our full model achieves the best scores under all the metrics. Without the coarse-level feature, the performance of *w/o f_c* drops significantly, revealing the importance of body prior. Without the double diffusion mechanism, *w/o DD* is hard to produce realistic novel views when the input is a sparse set of views. This manifests that our diffusion mechanism can effectively learn the body representation for the implicit field regression. Besides, most metrics are negatively changed with the absence of the transformer. This indicates that the transformer can effectively integrate multi-view information and produce more satisfactory results. The visual comparisons of the above variations are shown in Fig. 6.

5 Conclusion

We propose a novel double diffusion based neural radiance field to represent the human body implicitly from a sparse set of multi-view images. The proposed double diffusion mechanism fully exploits human body priors and image appearance details at both coarse and fine levels, thus leading to expressive representations of subject geometry and appearances. Benefiting from the diffused features, our proposed implicit field regression module can reconstruct high-quality human bodies. Thanks to our self-supervised training fashion, our method does not require any 3D supervision in training and is able to reconstruct human bodies of unseen subjects. Extensive experiments show that our method considerably outperforms concurrent works.

Acknowledgements

This work is partially supported by NSF China (62022072).

References

- [Chen *et al.*, 2021] Anpei Chen, Zexiang Xu, and Hao Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. *arXiv preprint arXiv:2103.15595*, 2021.
- [Dong and Fang, 2021] Junting Dong and Qi Fang. Fast and robust multi-person 3d pose estimation and tracking from multiple views. In *T-PAMI*, 2021.
- [Gilbert *et al.*, 2018] Andrew Gilbert, Marco Volino, John Collomosse, and Adrian Hilton. Volumetric performance capture from minimal camera viewpoints. In *ECCV*, pages 566–581, 2018.
- [Graham and Engelcke, 2018] Benjamin Graham and Engelcke. 3d semantic segmentation with submanifold sparse convolutional networks. In *CVPR*, pages 9224–9232, 2018.
- [Gropp and Yariv, 2020] Amos Gropp and Lior Yariv. Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020.
- [Kellnhofer *et al.*, 2021] Petr Kellnhofer, Lars C Jebe, and Gordon Wetzstein. Neural lumigraph rendering. In *CVPR*, pages 4287–4297, 2021.
- [Kingma and Ba, 2014] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [Kwon *et al.*, 2021] 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, 2021.
- [Loper *et al.*, 2015] Matthew Loper, Naureen Mahmood, and Javier Romero. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015.
- [Lorensen and Cline, 1987] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM siggraph*, 21(4):163–169, 1987.
- [Mildenhall *et al.*, 2020] Ben Mildenhall, Pratul P Srinivasan, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, pages 405–421. Springer, 2020.
- [Newell *et al.*, 2016] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, pages 483–499. Springer, 2016.
- [Niemeyer *et al.*, 2020] Michael Niemeyer, Lars Mescheder, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *CVPR*, pages 3504–3515, 2020.
- [Peng *et al.*, 2021] Sida Peng, Yuanqing Zhang, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *CVPR*, pages 9054–9063, 2021.
- [Saito *et al.*, 2019] Shunsuke Saito, Zeng Huang, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *CVPR*, pages 2304–2314, 2019.
- [Saito *et al.*, 2020] Shunsuke Saito, Tomas Simon, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *CVPR*, pages 84–93, 2020.
- [Shao *et al.*, 2021] Ruizhi Shao, Hongwen Zhang, Tao Yu, and Yebin Liu. Doublefield: Bridging the neural surface and radiance fields for high-fidelity human rendering. *arXiv preprint arXiv:2106.03798*, 2021.
- [Vaswani and Shazeer, 2017] Ashish Vaswani and Noam Shazeer. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017.
- [Wang *et al.*, 2021a] Peng Wang, Lingjie Liu, Yuan Liu, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021.
- [Wang *et al.*, 2021b] Qianqian Wang, Zhicheng Wang, and Kyle Genova. Ibrnet: Learning multi-view image-based rendering. In *CVPR*, 2021.
- [Yariv *et al.*, 2020] Lior Yariv, Yoni Kasten, and Dror Moran. Multiview neural surface reconstruction by disentangling geometry and appearance. *arXiv preprint arXiv:2003.09852*, 2020.
- [Yu *et al.*, 2018] Tao Yu, Zerong Zheng, and Yebin Liu. Doublefusion: Real-time capture of human performances with inner body shapes from a single depth sensor. In *CVPR*, pages 7287–7296, 2018.
- [Yu *et al.*, 2021] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *CVPR*, pages 4578–4587, 2021.
- [Zhang *et al.*, 2018] Richard Zhang, Phillip Isola, and Alexei Efros. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [Zheng *et al.*, 2021] Yang Zheng, Ruizhi Shao, Qionghai Dai, and Yebin Liu. Deepmulticap: Performance capture of multiple characters using sparse multiview cameras. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6239–6249, 2021.