

# Improving Neural Surface Reconstruction with Feature Priors from Multi-View Images

Xinlin Ren<sup>1\*</sup>, Chenjie Cao<sup>1,2\*</sup>, Yanwei Fu<sup>1†</sup>, and Xiangyang Xue<sup>1</sup>

<sup>1</sup> Fudan University

<sup>2</sup> Alibaba Group, DAMO Academy

{xlren20,cjcao20,yanweifu,xyxue}@fudan.edu.cn

**Abstract.** Recent advancements in Neural Surface Reconstruction (NSR) have significantly improved multi-view reconstruction when coupled with volume rendering. However, relying solely on photometric consistency in image space falls short of addressing complexities posed by real-world data, including occlusions and non-Lambertian surfaces. To tackle these challenges, we propose an investigation into feature-level consistent loss, aiming to harness valuable feature priors from diverse pretext visual tasks and overcome current limitations. It is crucial to note the existing gap in determining the most effective pretext visual task for enhancing NSR. In this study, we comprehensively explore multi-view feature priors from seven pretext visual tasks, comprising thirteen methods. Our main goal is to strengthen NSR training by considering a wide range of possibilities. Additionally, we examine the impact of varying feature resolutions and evaluate both pixel-wise and patch-wise consistent losses, providing insights into effective strategies for improving NSR performance. By incorporating pre-trained representations from MVSFormer [2] and QuadTree [38], our approach can generate variations of MVS-NeuS and Match-NeuS, respectively. Our results, analyzed on DTU and EPFL datasets, reveal that feature priors from image matching and multi-view stereo outperform other pretext tasks. Moreover, we discover that extending patch-wise photometric consistency to the feature level surpasses the performance of pixel-wise approaches. These findings underscore the effectiveness of these techniques in enhancing NSR outcomes. Codes and models are available at [https://github.com/maybeLx/MVS\\_NeuS](https://github.com/maybeLx/MVS_NeuS).

**Keywords:** Feature Priors · Neural Surface Reconstruction

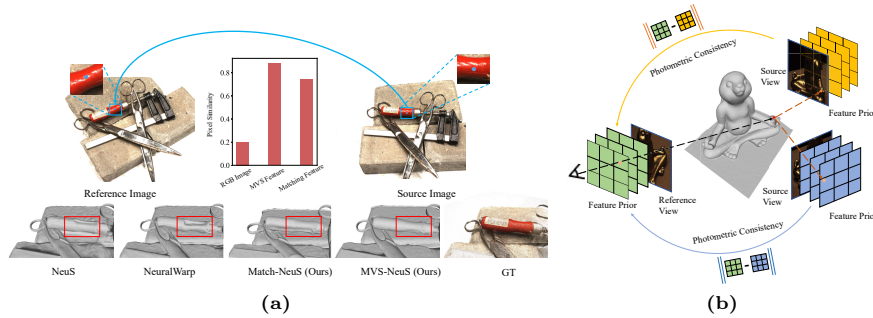
## 1 Introduction

Reconstructing surface from multi-view images is a crucial task in 3D computer vision. Traditionally, this task is addressed using a pipeline that includes depth prediction, point cloud reconstruction, and maximizing point consistency for surface mesh [1, 17, 21, 35]. Unfortunately, these traditional methods often result in

---

\*: Equal Contributions.

†: Corresponding Author



**Fig. 1:** (a) Pixel-wise similarity for RGB images: MVS features [2] and matching features [38] on non-Lambertian surfaces. (b) Multi-view consistency of feature priors: Given a ray direction, the geometry network can extract surface points, which are used to project and align source features to the reference view for optimizing patch-wise photometric consistency loss.

visual errors due to the complexity of these procedures. In contrast, recent approaches opt to represent the surface through volume rendering, inspired by the Neural Radiance Fields (NeRF) [28], thereby avoiding the accumulation of errors seen in the traditional pipeline. Notably, Neural Surface Reconstruction (NSR) has made significant strides in multi-view reconstruction, leveraging techniques such as Signed Distance Functions (SDF) [42, 51] and occupancy [31] combined with volume rendering. However, despite these advancements, these sophisticated methods still grapple with fundamental challenges in surface reconstruction, including maintaining geometric consistency across different views. This challenge becomes particularly pronounced when dealing with non-Lambertian surfaces and occlusions [6, 9], as illustrated in Fig. 1a.

To tackle these challenges, earlier studies [6, 9] have delved into optimizing patch-wise warping losses in image space, employing metrics such as Structural Similarity (SSIM) [45] or Normalized Cross-Correlation (NCC) [10]. In an effort to enhance the robustness of geometric consistency in real-world scenarios, Chen *et al.* [3] utilized features extracted from a learning-based Multi-View Stereo (MVS) model [55] to compute pixel-wise feature similarity. Despite promising results with MVS features, it remains unclear *whether the neural surface reconstruction can benefit from other pretext tasks and what key factors contribute to improved reconstruction*. Furthermore, it is worth investigating whether these features can also benefit from patch-wise photometric consistency.

This study aims to address these questions by conducting a comprehensive investigation into the impact of diverse pre-trained feature priors obtained from various pretext visual tasks for NSR learning. We thoroughly consider seven distinctly different pretext visual tasks, encompassing 2D supervised [4, 26, 47], 2D self-supervised [11, 15, 46], and 3D supervised [2, 13, 19, 24, 33, 38, 48] approaches. The pre-trained models from these tasks can serve as feature extractors for our NSR task. Crucially, we not only compare various architectures used in these models, ranging from Convolution Neural Networks (CNNs) [13, 24, 26, 48] to Vision Transformers (ViTs) [2, 15, 46, 47], but also explore feature maps from dif-

ferent layers within the same model to identify common factors relevant to NSR. Furthermore, we propose for extending patch-wise photometric consistency to the feature level, resulting in substantial improvements.

Formally, our approach enhances NSR by integrating various feature priors derived from multi-view images. Building upon the robust foundation of NeuS [42], which acts as the baseline method in this paper, we further improve NSR by introducing feature-level consistent loss from diverse pre-trained models, as depicted in Fig. 1b. Specifically, we compare 13 pre-trained models across 7 distinct pretext tasks, as listed in Tab. 1. To the best of our knowledge, this evaluation represents the most comprehensive assessment of pretext tasks on NSR to date, encompassing masked image modeling [11, 15, 46], image classification [4, 26], semantic segmentation [47], monocular depth estimation [19, 33], stereo matching [24, 48], image matching [38], and MVS [2, 13]. The DTU and EPFL datasets serve as the benchmark for evaluation [16].

Notably and interestingly, our empirical findings reveal that models pre-trained on MVS and image matching tasks consistently outperform others significantly. This success is attributed to the feature similarity supervision provided by depth annotation [16] for MVS and Structure from Motion (SfM) [23] for image matching. Moreover, these tasks yield high-resolution features, emerging as another crucial factor contributing to the superior performance of NSR models.

Furthermore, we observe that 3D-CNN regularization, designed to denoise the 3D cost volume [2, 13, 48], impedes effective 2D feature learning. Consequently, utilizing features from pre-trained models of image matching tasks without 3D cost volume regularization achieves the best performance and generalization in NSR. Moreover, in our experiments on the EPFL dataset [36], we make an intriguing observation that patch-wise photometric consistency can be effectively generalized to the feature level, resulting in remarkable improvements. Thus, bolstered by patch-wise photometric consistency losses from pre-trained models in both MVSFormer [2] and QuadTree [38] our approach can produce variations of MVS-NeuS and Match-NeuS. These variations demonstrate state-of-the-art performance on the DTU and EPFL datasets, respectively. Additionally, we take a further exploration of the feature priors on grid-based methods such as Neuralangelo [22], where feature priors still enjoy promising benefits. Thus our method could be extended to different neural representations (MLP or grid-based methods). Our contributions can be highlighted as follows:

- 1) We are the first to conduct a comprehensive evaluation of prior features derived from various pretext tasks on NSR. Additionally, we delve into a detailed comparison of pixel-wise feature consistent loss across these approaches. This comprehensive analysis advances our understanding of the effect of different feature priors for NSR.

- 2) Our research brings to light a significant finding: models pre-trained on image matching and MVS tasks consistently outperform others in the realm of NSR. This superior performance is attributed to the advantages gained from feature-similarity related pretext training and the utilization of high-resolution

**Table 1:** A summary of models from various pretext tasks including MIM (masked image modeling), IC (image classification), MDE (monocular depth estimation), SM (stereo matching), SS (semantic segmentation), IM (image matching), and MVS. Only the feature dimensions and resolutions of the highest and lowest features are shown. (Res: Resolution, Dim: Dimension)

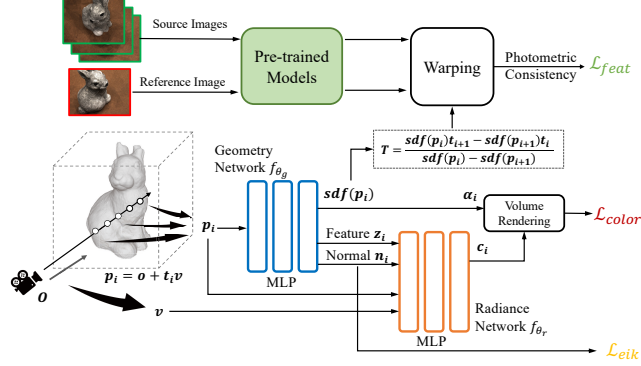
Task	Model	Architect	Feature Res	Feature Dim
MIM	MAE [15]	ViT	1/16	1024
	CovMAE [11]	ViT+CNN	1/16, 1/4	768, 256
	CroCo [46]	ViT	1/16	768
IC	Twins [4]	ViT	1/32, 1/4	1024, 128
	ConvNeXt [26]	CNN	1/32, 1/4	2048, 256
SS	SegFormer [47]	ViT	1/32, 1/4	512, 64
MDE	GLPN [19]	ViT	1/32, 1/4	512, 64
	MiDaS [33]	ViT	1/32	2048
SM	IGEV [48]	CNN	1/32, 1/4	160, 96
	RAFT-Stereo [24]	CNN	1/4	256
IM	QuadTree [38]	ViT+CNN	1/8, 1/2	256, 128
MVS	CascadeMVS [13]	CNN	1/8, 1	32, 8
	MVSFormer [2]	ViT+CNN	1/8, 1	32, 8

features. By identifying these key factors, our work contributes valuable insights that can inform future developments in NSR methods.

**3)** Another notable contribution of our research lies in the successful extension of the patch-wise photometric consistency loss to the feature level. This innovative approach stands as a significant advancement, leading to improved results regardless of different neural representations (MLP or grid-based methods). By bridging the gap between patch-wise and feature-level considerations, our work reaches the state-of-the-art in NSR.

## 2 Related Work

**Neural Surfaces Reconstruction (NSR).** Surface reconstruction based on deep learning has garnered significant attention due to its expressive representation capabilities. Representative algorithms focus on utilizing implicit representations to encode the geometry and appearance of 3D scenes through neural networks. These representations include occupancy fields [27, 31] and SDF [25, 49, 52, 56, 57]. Since the emergence of volume rendering techniques pioneered by NeRF [28], numerous NSR methods have seen a dramatic rise [3, 6, 9, 14, 32, 42–44, 51, 54]. Notably, VolSDF [51] and NeuS [42] convert volume density to SDF without the need for masks, while UNISURF [32] unifies both surface and volume rendering. Concurrently, grid-based representation methods are gaining attention [22, 29, 37, 41], as they offer faster convergence compared to implicit ones with competitive results. Essentially, note that these explicit NSR methods serve as robust foundations, and our contributions are orthogonal to these approaches. However, it is crucial to highlight that all these existing methods



**Fig. 2:** Approach overview. Our method is based on NeuS [42]. After achieving surface points from geometry network  $f_{\theta_g}$ , we further apply the multi-view consistent loss based on features from pre-trained models to improve the reconstruction quality.

overlook the importance of multi-view consistency during surface reconstruction. This gap forms the basis for our unique contributions in addressing this aspect. **Photometric Consistency.** Many previous studies have delved into the application of photometric consistency across various 3D reconstruction tasks [5, 18, 30, 34, 39, 40, 50, 53]. For SfM, several researches aim to optimize camera poses by minimizing photometric errors between corresponding pixels in different images [34, 39, 40]. Similarly, various self-supervised MVS methods [5, 18, 50] leverage consistent loss to supervise depth prediction. Moreover, in Simultaneous Localization and Mapping (SLAM), photometric consistency plays a crucial role. By comparing the appearance of projected 3D points in the image with the actual observed images, photometric consistency can detect potential errors or inaccuracies in geometry [30, 53].

**Photometric Consistency in NSR.** Prior works [6, 9, 14] have incorporated patch-wise photometric consistency directly computed from raw images or pixel-wise feature similarity to enhance NSR training. NeuralWarp [14] introduces a patch-wise warping loss along with a visibility map that masks incorrect warps arising from occlusion. Similarly, GeoNeus [9] and D-NeuS [6] focus on locating the zero-level set of geometry network to enforce multi-view geometry constraints, resulting in more accurate reconstruction outcomes. Despite these advancements, none of these methods thoroughly explore the distinctions between patch-wise and pixel-wise consistent loss, nor do they investigate the potential benefits of incorporating prior features from other pretext visual tasks as opposed to using raw images.

### 3 NSR based on Feature Priors

**Overview.** Given a set of posed multi-view images, the goal of NSR is to reconstruct high-quality surfaces by neural volume rendering, as illustrated in Fig. 2. We revisit the NeuS baseline [42] for multi-view consistency zero-crossing points.

We briefly introduce pretext tasks in Sec. 3.1 and discuss consistent loss in both pixel-wise and patch-wise terms in Sec. 3.2. By utilizing pre-trained representations from MVS and image matching pretext tasks, we have MVS-NeuS and Match-NeuS, which are compared in the experiments.

**NeuS** [42] proposes a novel scheme to achieve unbiased reconstructed surfaces in the first-order approximation of SDF. In this scheme, a ray  $\{p(t) = o + tv | t > 0\}$  is emitted from a camera center at  $o$  with view direction  $v$ . Colors  $\hat{C}(r)$  along a ray can be accumulated by volume rendering [8], which is approximated by  $N$  discretely sampled points as follows:

$$\hat{C}(r) = \sum_{i=1}^N T_i \alpha_i c_i, \quad T_i = \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (1)$$

where  $\alpha_i$  and  $T_i$  are the opacity and accumulated transmittance along the ray respectively. To achieve the unbiased surface, we can compute  $\alpha_i$  by :


$$\alpha_i = \max\left(\frac{\Phi_s(sdf(p_i)) - \Phi_s(sdf(p_{i+1}))}{\Phi_s(sdf(p_i))}, 0\right), \quad (2)$$

where  $sdf(\cdot)$  is the SDF function modeled by the geometry network  $f_{\theta_g}$ .  $\Phi_s(\cdot)$  is the sigmoid function and  $s$  is a learnable parameter.

**Differentiable Zero-Crossing Points.** Since achieving photometric consistency involves warping patches or pixels based on a specific depth, we define the depth planes through zero-crossing points achieved from the zero-level set of SDF function as [3, 6]. So it is crucial to apply photometric consistency specifically to the zero-crossing points of geometry network  $f_{\theta_g}$ , as shown in Fig. 3. We only require the first set of zero-crossing points  $T^*$  which is visible from the corresponding view. It can be expressed as:

$$T^* = \left\{ t | t = \frac{sdf(p_{t_i})t_{i+1} - sdf(p_{t_{i+1}})t_i}{sdf(p_{t_i}) - sdf(p_{t_{i+1}})}, t_i \in T \right\}, \quad (3)$$

$$T = \arg \min_i \{ t_i | sdf(p_{t_i}) \cdot sdf(p_{t_{i+1}}) < 0 \}. \quad (4)$$



**Fig. 3:** Two different ways to locate surface points during the network training at 5000 iterations. Compared with volume rendering, it is more accurate to find the surface points via Eq. 3. Accurate surface points are crucial for applying consistent loss.

### 3.1 Pretext Tasks

Various models in different pretext visual tasks have been explored as shown in Tab. 1. Our emphasis is specifically on the features obtained from the feature encoder modules of these pre-trained models. Importantly, many of these models include multi-scale feature maps. In this context, our focus is on the most representative features at both the highest and lowest scales, encompassing high-level semantic understanding and fine-grained low-level details, respectively.

**Masked Image Modeling (MIM)** has been widely investigated recently, as an effective self-supervised task [15]. MIM randomly masks input images at a relatively high rate and subsequently predicts the masked regions based on the context of unmasked areas. MIM has gained significant attention recently, with several pioneering works demonstrating its effectiveness in various downstream tasks [11, 15, 46]. In contrast to other self-supervised tasks, MIM exhibits favorable fine-grained feature correlations that can potentially benefit NSR learning.

**Image Classification (IC)** is a well-established supervised task known for extracting discriminative features. In this paper, we investigate models built on CNNs [26] and ViTs [4] that are pre-trained for image classification on ImageNet [7]. These models serve as baselines, helping us answer the question of whether engaging in such a classic pretext task also enhances the NSR.

**Semantic Segmentation (SS).** Models like [20, 47] have demonstrated outstanding performance in semantic segmentation by training on large-scale datasets. Unlike classification tasks, these models can capture more detailed, class-specific features, which holds the potential to enhance NSR training.

**Monocular Depth Estimation (MDE)** aims to predict the depth map of a scene using a single input view, essentially creating a 2D representation of a 3D scene. Several successful works [19, 33] have demonstrated the ability to recover fine details of the scene in this process.

**Stereo Matching (SM)** is the process of identifying corresponding pixels between two rectified images along epipolar lines, ultimately determining the disparity between these corresponding points. Noteworthy approaches like IGEV [48] and RAFT-Stereo [24] have demonstrated promising results in this task. The learning involved in establishing correspondences bears some similarities to the learning of consistent loss in NSR.

**Multi-View Stereo (MVS)** differs from stereo matching in that it involves multiple input views with known camera poses. However, the MVS problem often faces challenges such as occlusions and varying light conditions. Consequently, MVS methods like CasMVSNet [13] and MVSFormer [52] need to enhance feature learning to accurately identify corresponding pixels along epipolar lines. It's worth noting that MVS methods may encounter data bias, as discussed in Sec. 4.2, particularly on the DTU dataset [16].

**Image Matching (IM).** In comparison to stereo matching and MVS, image matching tackles a more intractable problem, *i.e.*, detecting fine-grained matching pairs between two images without assistance from known camera poses. To solve this challenge, Tand et al. [38] utilize a transformer with a lightweight QuadTree attention mechanism to capture long-range dependencies. Despite

working without constraints from camera poses, we observe that image matching methods can still achieve effective feature representation.

### 3.2 Feature-Level Photometric Consistency

To leverage the consistency of geometry across different views, we utilize the features  $\mathbf{F}_r$  and  $\mathbf{F}_s$  extracted from the reference and source images through various pre-trained models respectively. Photometric consistency can be applied to these multi-view features to enforce their alignment and coherence. Specifically, from the reference viewpoint, we locate a surface point  $p$  derived from Eq. 3 with its 2D projection coordinate denoted as  $x$ . This 2D point is then associated with the corresponding point  $x'$  in the source images through the homography  $H$ :

$$x' = Hx, \quad n^T p + d = 0, \quad (5)$$

$$H = K_s R_s (I - \frac{(R_s^{-1} t_s - R_r^{-1} t_r) n^T R_r}{d}) R_r^{-1} K_r^{-1}, \quad (6)$$

where  $K_s$  and  $K_r$  denote the internal calibration matrix of source and reference cameras,  $(R_s, R_r)$  and  $(t_s, t_r)$  are the  $3 \times 3$  rotation matrix, and 3D translation vector of two views, respectively.  $n$  denotes the surface normal, which can be computed by automatic differentiation of the geometry network at each point  $p$ . Additionally,  $I$  represents the identity matrix and  $d$  is the distance from the camera center to the plane of the zero-crossing point (Eq. 3).

**Pixel Warping.** To compute the pixel-level consistency loss, we can evaluate the feature difference between the pixel  $x$  in the reference view and its corresponding pixel  $x'$  in the source views:

$$\mathcal{L}_{pixel} = D_{pixel}(\mathbf{F}_r(x), \mathbf{F}_s(x')), \quad (7)$$

where  $D_{pixel}$  is the cosine similarity function to measure the feature similarity in pixel-wise.

**Patch Warping.** To more effectively incorporate local feature information, we can broaden the concept of pixel warping to patch warping. To define a small patch  $s$  on the surface, traditionally we would compute the surface points using Eq. 3 and sample numerous points around those surface points to construct the patch. However, this is computationally expensive. As an alternative, we use the central point  $\hat{s}$  of the patch  $s$  to represent it and derive the homography matrix  $H$ . We can perform patch warping from the source view to the reference view and measure the difference between them. The patch-wise consistent loss is

$$\mathcal{L}_{patch} = D_{patch}(\mathbf{F}_r(\hat{s}), \mathbf{F}_s(\hat{s}')), \quad (8)$$

where  $\hat{s}$ ,  $\hat{s}'$  denote the centers of reference view patches and projected source view patches respectively.  $D_{patch}$  can be represented as NCC [10], SSIM [45], or patch similarity which measures the average of similarity between corresponding elements in the two patches. Since grid-based manners [22] can also locate the



surface points via Eq. 3, we easily extend our method to the grid-based methods in Sec. 4.3. More detailed formulas will be explained in the supplementary.

**About Occlusion.** Although we can utilize photometric consistency to supervise the geometry network and encourage accurate surface learning, it’s important to note that not all surface points are visible in source view images. Choosing unsuitable source images would cause prominent degradation to the reconstruction. To address this issue, we employ a simple yet effective strategy called Top-k view selection. Specifically, we sort and select 10 source views based on their relative angles from small to large to the reference perspective. Subsequently, we choose Top-k views to compute the patch-wise consistent loss according to the best k photometric scores. We empirically set k to 4 with proper generalization for most scenes. However, we have observed that this strategy does not provide significant benefits to pixel-wise consistency due to the incomplete local feature information which results in unstable view selection during training (Sec. 4.2).

**Training Loss.** To render colors from a specific view, the overall loss function is defined as:

$$\mathcal{L} = \mathcal{L}_{color} + \lambda_1 \cdot \mathcal{L}_{eik} + \lambda_2 \cdot \mathcal{L}_{feat}, \quad (9)$$

where  $\mathcal{L}_{color} = \sum_{r \in \mathcal{R}} \|\hat{C}(\mathbf{r}) - C(\mathbf{r})\|_1$  is the  $\ell_1$  loss from minibatch of rays  $\mathcal{R}$ .  $\mathcal{L}_{eik}$  is the eikonal loss defined in [12] to regularize the SDF values predicted by the geometry network  $f_{\theta_g}$  as depicted in Fig. 2.

We use feature loss, denoted as  $\mathcal{L}_{feat}$ , and it is implemented as  $\mathcal{L}_{patch}$ . Additionally, in Sec. 4.1, we apply the feature loss as  $\mathcal{L}_{pixel}$  to assess the features of pre-trained models derived from different pretext tasks. Our results demonstrate that  $\mathcal{L}_{patch}$  exhibits greater robustness compared to  $\mathcal{L}_{pixel}$ , emphasizing the effectiveness of our patch warping loss. During the training stage, we set  $\lambda_1$  and  $\lambda_2$  as 0.1 and 0.5, respectively.

**Discussion.** While NeuralWarp [14] has introduced a patching-based warping method to transform source view patches to the reference view, it posits that every sampled point along the ray could be projected to other views to obtain the warped colors, leading to the final warped colors through volume rendering. In contrast, our approach prioritizes the surface points, offering a more accurate and reasonable method that does not depend on visible maps. Furthermore, we extend the patch warping from RGB pixel-level to feature-level, significantly improving the overall quality of reconstruction.

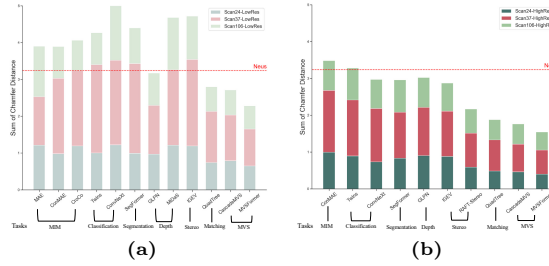
## 4 Experiments

**Datasets.** The DTU dataset tested in NSR comprises 15 scenes with 49 or 64 multi-view images collected in an indoor environment with fixed camera poses following [52]. The Chamfer Distance (CD) is utilized to evaluate the performance of reconstructed results on the DTU. We also test two outdoor scenes in EPFL dataset [36], which has 7 and 11 images at a resolution of  $2048 \times 3072$ , and one of the scenes is particularly challenging due to the occlusions caused by limited viewpoints.

**Implementation Details.** We use the same architecture as [42]. During the training, 512 rays are sampled for each batch following the hierarchical sampling strategy in NeuS [42]. The learning rate is linearly warmed up from 0 to  $5e-3$  in the first 5k steps and then decayed to  $2.5e-5$  with a cosine schedule. For each scene, we train the network for 300k steps costing 10 hours on a single A6000.

#### 4.1 Experiments for Different Pretext Tasks

**Settings.** We evaluate the effect of various pretext tasks in Tab. 1 by applying pixel similarity as the consistency loss on a subset of the DTU dataset including Scan24, Scan37, and Scan106. These objects are relatively challenging due to difficult lighting conditions and occlusions. Because of the quadratic memory footprint of ViT in high-resolution features, we downsample all images into  $512 \times 768$  for a fair comparison. We only consider the highest and lowest features from the feature encoder as mentioned in Sec. 3.1. More experiments on multi-scale features are presented in the supplementary.



**Fig. 4:** Quantitative results of using prior features from different pretext tasks for two scenarios: (a) with the lowest resolution features (high-level features), and (b) with the highest resolution features (low-level features). The red lines indicate the performance of NeuS, with lower values indicating better performance in terms of Chamfer Distance. High-resolution features generally outperform the low-resolution ones.

**In-depth Discussions.** As depicted in Fig. 4, when employing low-resolution features for the same model, the performance is noticeably worse compared to using high-resolution features. This is because projecting surface points onto low-resolution feature maps does not yield fine-grained effective features, thus hindering the geometry network from accurately learning the correct geometry. Furthermore, models optimized with high-resolution features from MIM, image classification, semantic segmentation, and monocular depth estimation only perform slightly better than the baseline NeuS [42]. For classification [4, 26] and segmentation [47], we think that these pretext tasks focus on high-level semantic information that is not suitable for NSR. For MIM [11, 15, 46], the most concern is that the effective feature correlations are not fine-grained enough, which limits the performance. For the depth estimation task [19, 33], we think that it shares little connection with NSR learning.

In contrast, models being optimized with feature priors from stereo matching, image matching, and MVS achieve good performance. These tasks share an


**Table 2:** Quantitative results of CD (mm) on the EPFL dataset with different consistent losses. (Fou: Fountain, Her: Herzjesu)

Model	Pre-Trained	Consistent Loss	Fou	Her	Mean
Match-NeuS	QuadTree [38]	Pixel Similarity	1.288	1.832	1.560
		Patch Similarity	1.388	1.722	1.555
		Patch NCC	<b>1.206</b>	<b>1.622</b>	<b>1.414</b>
		Patch SSIM	1.218	1.678	1.448
MVS-NeuS	MVSFormer [2]	Pixel Similarity	<b>1.199</b>	1.698	1.449
		Patch Similarity	<u>1.212</u>	1.689	1.451
		Patch NCC	1.242	<u>1.646</u>	<u>1.444</u>
		Patch SSIM	1.261	1.706	1.484
NeuralWarp [42]	/	Patch SSIM	1.92	2.03	1.97
NeuS [42]	/	/	1.741	5.613	3.677

important common factor, *i.e.*, matching feature similarity across different view-points, while the supervision is usually annotated by the depth captured from SfM [23] or specific devices [16]. Thus, low-level and fine-grained feature correlations should be considered in these manners. In general, using feature priors from the same task generally results in a similar performance. However, in the stereo matching task, there is a significant performance gap between IGEV [48] and RAFT-Stereo [24], even though these two models have a similar architecture except for the cost volume regularization. Specifically, IGEV builds the cost volume as a regularization module that encodes and denoises feature conflicts caused by occlusions and non-Lambertian in the image level through 3D CNNs. Although 3D cost volume addresses the issues from the 2D feature encoder, we clarify that it could not help to learn essentially better 2D feature representation. As evidence, features from IGEV achieve inferior performance than RAFT-Stereo, because of the lightweight local cost volume GRU in RAFT-Stereo [24], which releases the capacity of 2D features learned by the feature encoder. From Fig. 4, we observe that QuadTree [38] pre-trained with image matching and MVS methods [2, 13] enhanced NeuS achieve much better results compared with others. As mentioned in Tab. 1, models pre-trained on these two tasks enjoy more fine-grained features (1/2 and 1:1). MVS methods are all pre-trained on the DTU dataset [16], while the matching method is pre-trained on the MegaDepth [23]. So MVS manners may suffer from a little data leakage in Fig. 4. To be fair, we further compare the performance of image matching and MVS on EPFL dataset in Sec. 4.2.


## 4.2 MVS and Image Matching based NSR

By incorporating features of pre-trained models such as MVSFormer [2] and QuadTree [38], our approach can generate variants of MVS-NeuS and Match-NeuS, respectively. Additionally, the evaluation is conducted on EPFL [36] and includes both pixel-wise and patch-wise losses.



**Fig. 5:** Qualitative results of different source view selecting strategies based on pixel similarity and patch similarity consistent losses in the Herzjesu of EPFL dataset.

From Tab. 2, we find that the patch-wise NCC and SSIM could be successfully extended to feature level, and achieve better results compared with pixel similarity for Match-NeuS. Additionally, Match-NeuS with patch NCC outperforms MVS-NeuS, despite features from MVSFormer containing larger resolutions. As discussed in Sec. 4.1, the main reason for the better performance of Match-NeuS is that image matching models are pre-trained without the cost volume regularization, because no camera poses could be used to formulate the 3D cost volume. Furthermore, due to the prohibitive computation for the 3D CNN in the cost volume regularization, MVS methods should reduce the 2D feature channels for efficiency. Thus the matching-based QuadTree has 256 feature channels, while MVSformer only contains 8 feature channels with less capacity. In Fig. 6, we show that Match-NeuS with patch NCC recovers more detailed structures in Fountain, while reconstructing smoother walls.



**Fig. 6:** Qualitative comparisons on EPFL dataset.

**Table 3:** Ablation study of CD (mm) on matching features **w/wo** attention modules. We further compare them to the results of high-resolution features (1/2) without attention modules.


Model	Consistent Loss Res.	Fountain	Herzjesu	Mean
Match-NeuS	Pixel Similarity 1/2	<b>1.288</b>	<b>1.832</b>	<b>1.560</b>
Match-NeuS	Pixel Similarity 1/8	1.634	3.360	2.497
+ Attention	Pixel Similarity 1/8	1.502	3.010	2.256

**View Selection.** We compare the view selection for both pixel-wise and patch-wise losses in Fig. 5, From Fig. 5, the Top-k view selection mentioned in Sec. 3.2 fails to adaptively choose reasonable source views through pixel-wise feature similarity, resulting in poor reconstruction. Thus we need to manually select nearby source views to stabilize the training of pixel-wise feature similarity. However, the Top-k view selection based on patch-wise feature similarity can

achieve stable training and better reconstruction. Because local patches usually capture more information with larger perceptive fields, effectively addressing occlusion issues.

**Features with Attention Modules.** We observe that QuadTree employs the interleaved attention mechanism [38] including self and cross-attention modules. Attention modules enable explicit feature matching across different views and enhance the feature capacity. However, due to its high computational cost, the transformer layer is designed to refine only  $1/8$  feature scale. Hence this improvement fails to be transferred to high-resolution feature learning ( $1/2$ ). We have presented a comparison in Tab. 3 to showcase this improvement. It is evident that by incorporating the features after attention learning, Match-NeuS improves its performance, but it still performs worse than high-resolution feature learning.

#### 4.3 Extension of Feature Prior on Grid-based Representation.




**Fig. 7:** Qualitative results of Neuralangelo with feature priors.

We incorporate feature priors into grid-based Neuralangelo [22] as in Fig 7. Specifically, we train Neuralangelo with MVSFormer feature priors on DTU dataset for 300k iterations. The inclusion of MVS features enhances the accuracy of geometry and accelerates convergence, especially in intricate objects like scissors’ handles in Scan37 and the snowman in Scan69. Moreover, Neuralangelo with feature priors achieves smoother flat surfaces in Scan40 and Scan63.

#### 4.4 Full Models Compared on DTU

We compare MVS-NeuS and Match-NeuS with patch NCC as the consistent loss to other competitors on 15 scenes in the DTU dataset. Tab. 4 shows the quantitative results. Notably, our methods outperform other competitors by a large margin. Furthermore, we observed the qualitative results in Fig. 8 that our method can accurately reconstruct the geometry of these scenes. Specifically, our method is capable of not only modeling smooth flat planes but also clearly capturing gaps between two planes or the handle of scissors. In contrast, many other methods tend to fuse together in Scan37 with obvious artifacts. Quantitative results in Tab. 4 show that our model significantly outperforms other competing approaches.



**Fig. 8:** Comparisons of surface reconstruction in DTU dataset.

**Table 4:** Results of CD (mm) compared on DTU dataset (lower is better). The upper part methods enjoy additional annotations such as masks (IDR and COLMAP) or custom filtering (MVSDF). Note that the results of GeoNeuS\* are re-trained by ourselves, The best results are in **bold**.

Scan	24	37	40	55	63	65	69	83	97	105	106	110	114	118	122	Mean
COLMAP [35]	0.45	0.91	0.37	0.37	0.90	1.00	0.54	1.22	1.08	0.64	0.48	0.59	0.32	0.45	0.43	0.65
MVSDF [56]	0.93	1.75	0.88	0.44	1.11	0.90	0.75	1.26	1.02	1.35	0.87	0.84	0.34	0.47	0.47	0.88
IDR [52]	1.63	1.87	0.63	0.48	1.04	0.79	0.77	1.33	1.16	0.76	0.67	0.90	0.42	0.51	0.53	0.90
NeRF [28]	1.90	1.60	1.85	0.58	2.28	1.27	1.47	1.67	2.05	1.07	0.88	2.53	1.06	1.15	0.96	1.49
NeuS [42]	1.37	1.21	0.73	0.40	1.20	0.70	0.72	<b>1.01</b>	1.16	0.82	0.66	1.69	0.39	0.49	0.51	0.87
VolSDF [51]	1.14	1.26	0.81	0.49	1.25	0.70	0.72	1.29	1.18	0.70	0.66	1.08	0.42	0.61	0.55	0.86
NeuralWarp [6]	0.49	0.71	0.38	0.38	<b>0.79</b>	0.81	0.82	1.20	1.06	0.68	0.66	0.74	0.41	0.63	0.51	0.68
GeoNeuS* [9]	0.40	0.69	0.36	0.37	<b>0.79</b>	0.78	<b>0.48</b>	1.17	0.88	<b>0.60</b>	0.50	0.73	0.30	0.39	0.41	0.59
D-NeuS [3]	0.44	0.79	0.35	0.39	0.88	0.58	0.55	1.35	0.91	0.76	<b>0.40</b>	0.72	0.31	0.39	<b>0.39</b>	0.61
PET-NeuS [44]	0.56	0.75	0.68	0.36	0.87	0.76	0.69	1.33	1.08	0.66	0.51	1.04	0.34	0.51	0.48	0.71
MVS-NeuS	<b>0.37</b>	<b>0.63</b>	<b>0.31</b>	<b>0.32</b>	<b>0.79</b>	<b>0.51</b>	0.50	1.12	<b>0.87</b>	0.62	0.46	<b>0.64</b>	<b>0.29</b>	<b>0.38</b>	0.41	<b>0.548</b>
Match-NeuS	0.39	0.67	0.32	0.33	<b>0.79</b>	<b>0.51</b>	0.49	1.21	<b>0.87</b>	0.63	0.48	0.71	<b>0.29</b>	<b>0.38</b>	0.41	0.565

## 5 Conclusion

In this paper, we conducted an evaluation of various pretext tasks and consistent losses for NSR. Our findings showed that the feature priors from some pretext visual tasks including MVS and image matching can significantly improve the NSR performance. We also empirically observed that features with higher resolution consistently outperform the lower ones. Additionally, we further explored optimizing MVS and image matching enhanced NSR with both pixel-wise and patch-wise consistent losses. The results showed that the patch-wise consistent loss could be successfully generalized to the feature level and achieve superior reconstruction for the matching-based Match-NeuS.

## Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant 62076067 and Shanghai Municipal Science and Technology Major Project (NO. 2021SHZDZX0103). This work was performed using the CFFF platform of Fudan University. Yanwei Fu is with School of Data Science, Fudan ISTBI—ZJNU Algorithm Centre for Brain-inspired Intelligence, Fudan University, Shanghai Key Lab of Intelligent Information Processing, and Technology Innovation Center of Calligraphy and Painting Digital Generation, Ministry of Culture and Tourism, China. Email: yanweifu@fudan.edu.cn

## References

1. Barnes, C., Shechtman, E., Finkelstein, A., Goldman, D.B.: Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.* **28**(3), 24 (2009)
2. Cao, C., Ren, X., Fu, Y.: Mvsformer: Multi-view stereo by learning robust image features and temperature-based depth. *Transactions of Machine Learning Research*
3. Chen, D., Zhang, P., Feldmann, I., Schreer, O., Eisert, P.: Recovering fine details for neural implicit surface reconstruction. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 4330–4339 (2023)
4. Chu, X., Tian, Z., Wang, Y., Zhang, B., Ren, H., Wei, X., Xia, H., Shen, C.: Twins: Revisiting the design of spatial attention in vision transformers. *Advances in Neural Information Processing Systems* **34**, 9355–9366 (2021)
5. Dai, Y., Zhu, Z., Rao, Z., Li, B.: Mvs2: Deep unsupervised multi-view stereo with multi-view symmetry. In: *2019 International Conference on 3D Vision (3DV)*. pp. 1–8. Ieee (2019)
6. Darmon, F., Bascle, B., Devaux, J.C., Monasse, P., Aubry, M.: Improving neural implicit surfaces geometry with patch warping. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6260–6269 (2022)
7. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *2009 IEEE conference on computer vision and pattern recognition*. pp. 248–255. Ieee (2009)
8. Drebin, R.A., Carpenter, L., Hanrahan, P.: Volume rendering. *ACM Siggraph Computer Graphics* **22**(4), 65–74 (1988)
9. Fu, Q., Xu, Q., Ong, Y.S., Tao, W.: Geo-neus: geometry-consistent neural implicit surfaces learning for multi-view reconstruction. *arXiv preprint arXiv:2205.15848* (2022)
10. Galliani, S., Lasinger, K., Schindler, K.: Massively parallel multiview stereopsis by surface normal diffusion. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 873–881 (2015)
11. Gao, P., Ma, T., Li, H., Dai, J., Qiao, Y.: Convmae: Masked convolution meets masked autoencoders. *arXiv preprint arXiv:2205.03892* (2022)
12. Gropp, A., Yariv, L., Haim, N., Atzmon, M., Lipman, Y.: Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099* (2020)
13. Gu, X., Fan, Z., Zhu, S., Dai, Z., Tan, F., Tan, P.: Cascade cost volume for high-resolution multi-view stereo and stereo matching. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 2495–2504 (2020)

14. Guo, H., Peng, S., Lin, H., Wang, Q., Zhang, G., Bao, H., Zhou, X.: Neural 3d scene reconstruction with the manhattan-world assumption. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5511–5520 (2022)
15. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16000–16009 (2022)
16. Jensen, R., Dahl, A., Vogiatzis, G., Tola, E., Aanaes, H.: Large scale multi-view stereopsis evaluation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 406–413 (2014)
17. Kazhdan, M., Hoppe, H.: Screened poisson surface reconstruction. *ACM Transactions on Graphics (ToG)* **32**(3), 1–13 (2013)
18. Khot, T., Agrawal, S., Tulsiani, S., Mertz, C., Lucey, S., Hebert, M.: Learning unsupervised multi-view stereopsis via robust photometric consistency. *arXiv preprint arXiv:1905.02706* (2019)
19. Kim, D., Ka, W., Ahn, P., Joo, D., Chun, S., Kim, J.: Global-local path networks for monocular depth estimation with vertical cutdepth. *arXiv preprint arXiv:2201.07436* (2022)
20. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. *arXiv preprint arXiv:2304.02643* (2023)
21. Labatut, P., Pons, J.P., Keriven, R.: Efficient multi-view reconstruction of large-scale scenes using interest points, delaunay triangulation and graph cuts. In: 2007 IEEE 11th international conference on computer vision. pp. 1–8. IEEE (2007)
22. Li, Z., Müller, T., Evans, A., Taylor, R.H., Unberath, M., Liu, M.Y., Lin, C.H.: Neuralangelo: High-fidelity neural surface reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8456–8465 (2023)
23. Li, Z., Snavely, N.: Megadepth: Learning single-view depth prediction from internet photos. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2041–2050 (2018)
24. Lipson, L., Teed, Z., Deng, J.: Raft-stereo: Multilevel recurrent field transforms for stereo matching. In: 2021 International Conference on 3D Vision (3DV). pp. 218–227. IEEE (2021)
25. Liu, S., Zhang, Y., Peng, S., Shi, B., Pollefeys, M., Cui, Z.: Dist: Rendering deep implicit signed distance function with differentiable sphere tracing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2019–2028 (2020)
26. Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S.: A convnet for the 2020s. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022)
27. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4460–4470 (2019)
28. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
29. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)* **41**(4), 1–15 (2022)



30. Mur-Artal, R., Tardós, J.D.: Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras. *IEEE transactions on robotics* **33**(5), 1255–1262 (2017)
31. Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3504–3515 (2020)
32. Oechsle, M., Peng, S., Geiger, A.: Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5589–5599 (2021)
33. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence* **44**(3), 1623–1637 (2020)
34. Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4104–4113 (2016)
35. Schönberger, J.L., Zheng, E., Frahm, J.M., Pollefeys, M.: Pixelwise view selection for unstructured multi-view stereo. In: *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*. pp. 501–518. Springer (2016)
36. Strecha, C., Von Hansen, W., Van Gool, L., Fua, P., Thoennessen, U.: On benchmarking camera calibration and multi-view stereo for high resolution imagery. In: *2008 IEEE conference on computer vision and pattern recognition*. pp. 1–8. Ieee (2008)
37. Sun, C., Sun, M., Chen, H.T.: Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5459–5469 (2022)
38. Tang, S., Zhang, J., Zhu, S., Tan, P.: Quadtree attention for vision transformers. *ICLR* (2022)
39. Teed, Z., Deng, J.: Deepv2d: Video to depth with differentiable structure from motion. *arXiv preprint arXiv:1812.04605* (2018)
40. Vijayanarasimhan, S., Ricco, S., Schmid, C., Sukthankar, R., Fragkiadaki, K.: Sfmnet: Learning of structure and motion from video. *arXiv preprint arXiv:1704.07804* (2017)
41. Wang, J., Bleja, T., Agapito, L.: Go-surf: Neural feature grid optimization for fast, high-fidelity rgb-d surface reconstruction. *arXiv preprint arXiv:2206.14735* (2022)
42. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689* (2021)
43. Wang, Y., Skorokhodov, I., Wonka, P.: Hf-neus: Improved surface reconstruction using high-frequency details. *Advances in Neural Information Processing Systems* **35**, 1966–1978 (2022)
44. Wang, Y., Skorokhodov, I., Wonka, P.: Pet-neus: Positional encoding tri-planes for neural surfaces. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12598–12607 (2023)
45. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004)

46. Weinzaepfel, P., Leroy, V., Lucas, T., Brégier, R., Cabon, Y., Arora, V., Antsfeld, L., Chidlovskii, B., Csurka, G., Revaud, J.: Croco: Self-supervised pre-training for 3d vision tasks by cross-view completion. *arXiv preprint arXiv:2210.10716* (2022)
47. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems* **34**, 12077–12090 (2021)
48. Xu, G., Wang, X., Ding, X., Yang, X.: Iterative geometry encoding volume for stereo matching. *arXiv preprint arXiv:2303.06615* (2023)
49. Xu, Q., Wang, W., Ceylan, D., Mech, R., Neumann, U.: Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. *Advances in neural information processing systems* **32** (2019)
50. Yang, Z., Wang, P., Xu, W., Zhao, L., Nevatia, R.: Unsupervised learning of geometry from videos with edge-aware depth-normal consistency. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 32 (2018)
51. Yariv, L., Gu, J., Kasten, Y., Lipman, Y.: Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems* **34**, 4805–4815 (2021)
52. Yariv, L., Kasten, Y., Moran, D., Galun, M., Atzmon, M., Ronen, B., Lipman, Y.: Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems* **33**, 2492–2502 (2020)
53. Yi, X., Zhou, Y., Habermann, M., Golyanik, V., Pan, S., Theobalt, C., Xu, F.: EgoLocate: Real-time motion capture, localization, and mapping with sparse body-mounted sensors. *arXiv preprint arXiv:2305.01599* (2023)
54. Yu, Z., Peng, S., Niemeyer, M., Sattler, T., Geiger, A.: Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *arXiv preprint arXiv:2206.00665* (2022)
55. Zhang, J., Yao, Y., Li, S., Luo, Z., Fang, T.: Visibility-aware multi-view stereo network. *arXiv preprint arXiv:2008.07928* (2020)
56. Zhang, J., Yao, Y., Quan, L.: Learning signed distance field for multi-view surface reconstruction. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 6525–6534 (2021)
57. Zhang, K., Luan, F., Wang, Q., Bala, K., Snavely, N.: 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*. pp. 5453–5462 (2021)