RC-MVSNet: Unsupervised Multi-View Stereo with Neural Rendering

Di Chang¹, Aljaž Božič¹, Tong Zhang², Qingsong Yan³, Yingcong Chen³, Sabine Süsstrunk², Matthias Nießner¹

 ¹ Technical University of Munich
 ² École Polytechnique Fédérale de Lausanne
 ³ Hong Kong University of Science and Technology di.chang@tum.de

Abstract. Finding accurate correspondences among different views is the Achilles' heel of unsupervised Multi-View Stereo (MVS). Existing methods are built upon the assumption that corresponding pixels share similar photometric features. However, multi-view images in real scenarios observe non-Lambertian surfaces and experience occlusions. In this work, we propose a novel approach with neural rendering (RC-MVSNet) to solve such ambiguity issues of correspondences among views. Specifically, we impose a depth rendering consistency loss to constrain the geometry features close to the object surface to alleviate occlusions. Concurrently, we introduce a reference view synthesis loss to generate consistent supervision, even for non-Lambertian surfaces. Extensive experiments on DTU and Tanks&Temples benchmarks demonstrate that our RC-MVSNet approach achieves state-of-the-art performance over unsupervised MVS frameworks and competitive performance to many supervised methods. The code is released at https://github.com/Boese0601/ RC-MVSNet.

Keywords: End-to-end Unsupervised Multi-View Stereo, Neural Rendering, Depth Estimation

1 Introduction

Multi-View Stereo (MVS) [19] is a long-standing and fundamental task in 3D computer vision. MVS aims to recover a 3D point cloud of real scenes from multi-view images and corresponding calibrated cameras. The widespread application of deep learning in recent years has lead to the emergence of end-toend MVS depth estimation networks. A popular learning-based pipeline MVS-Net [33] proposes to encode RGB information from different camera views into a cost volume and predict depth maps for point cloud reconstruction. Follow-up fully-supervised methods [8,14,21,29] further improved the neural network architecture, lowering memory usage and achieving state-of-the-art depth estimation performance on several benchmarks [2,12]. However, these methods heavily rely



Fig. 1. Qualitative comparison of 3D reconstruction of our RC-MVSNet and previous unsupervised methods. [11,26]

on ground truth depth for supervision, which requires a depth sensor to collect the training data. It restricts these methods to limited datasets and largely indoor settings. To make MVS practical in more general real-world scenarios, it is vital to consider alternative unsupervised learning-based methods that can provide competitive accuracy compared to the supervised ones, while not requiring any ground truth depth.



Fig. 2. (left) Real-world photometric conditions and occlusions. (right) Top: Failure cases of a photometric consistency assumption. Bottom: The results of our proposed rendering consistency approach.

Existing unsupervised MVS methods [6,9,11,26] are based on a *photometric* consistency hypothesis, which states that pixels belonging to the same 3D point have identical color properties under different view directions. However, in a real-world environment, as shown in Fig. 2, occlusions, reflecting, non-Lambertian

3

surfaces, varying camera exposure, and other variables will make such a hypothesis invalid. Thus, inaccurately matched points lead to false geometric reconstruction. Although Xu et al. [26] proposes to use semantic clues to remove the ambiguity in stereo matching, the improvement is very marginal due to its high dependency on the accuracy of a pretrained semantic features extraction model (with supervision). To remove inconsistencies with supervision, in this paper, we propose to use neural rendering to solve ambiguity in the case of view-dependent photometric effects and occlusions.

Recently, there has been an increasing interest in novel view synthesis, particularly with the introduction of Neural Radiance Fields [15] that can model view-dependent photometric effects based on differentiable volumetric rendering. While initially focused on per-scene optimization with densely sampled input views, follow-up approaches [4,18,22] propose to use a 2D CNN encoder that can be trained to predict novel views even with very few input images. Aside from the view-dependent radiance that defines the color along the viewing ray, these methods also learn a volume density, which can be interpreted as depth when integrated over the viewing ray. It is worth noting that the depth is learned in a purely unsupervised fashion. However, since the main goal is novel view synthesis, the depth obtained from volume density is often inaccurate. In our method, we build from the volumetric rendering approach proposed in [15] with novel loss functions that can resolve view-dependent effects and occlusions, while keeping the depth representation used in MVS approaches to ensure locally accurate and smooth depth predictions.

To this end, we introduce RC-MVSNet, a novel end-to-end differentiable network for *unsupervised multi-view stereo*. Combining the benefits of viewdependent rendering and structured depth representation, our method achieves state-of-the-art depth prediction results in the competitive DTU benchmark [2], and also demonstrates robust performance on out-of-distribution samples in the Tanks and Temples dataset [12]. In summary, our contributions are the following:

- We propose a reference view synthesis loss based on neural volumetric rendering to generate RGB supervision that is able to account for view-dependent photometric effects.
- We introduce Gaussian-Uniform mixture sampling to learn the geometry features close to the object surface to overcome occlusion artefacts present in existing unsupervised MVS approaches.
- We introduce a depth rendering consistency loss to refine the initial depth map by depth priors and ensure the robustness and smoothness of the prediction.

2 Related Work

2.1 Supervised Multi-View Stereo

Many *supervised* approaches have proposed to use CNNs to interpret 3D scenes by predicting depth maps of the RGB inputs and reconstructing the point cloud

using depth map filtering. Most state-of-the-art methods employ 3D cost volumes. As a representative work, MVSNet [33] encodes camera parameters and features into cost volumes through homography warping and regularizes the volume by 3D CNNs to generate depth prediction. Following works, e.g., [8,21,32] improve the performance of MVSNet and reduce the memory cost by introducing a multi-stage architecture and learn depth prediction in a coarse-to-fine manner. Furthermore, [25,30,34] replace dense 3D convolution with convolutional recurrent GRU or LSTM units. However, the reliance on ground truth depth limits their application to specific datasets as discussed in Sec. 1. Therefore, it is necessary to explore alternative unsupervised methods.

2.2 Unsupervised and Self-supervised Multi-View Stereo

End-to-end unsupervised and pseudo-label-based multi-stage self-supervised learning play pivotal roles in 3D vision, especially in multi-view reconstruction systems. The fundamental assumption of photometric consistency provides feasibility for unsupervised MVS. For instance, Unsup_MVS [11] proposed the first end-to-end learning-based unsupervised MVS framework: the images of source views are inversely warped to the reference view with its predicted depth map, and photometric consistency and SSIM [23] are enforced to minimize the discrepancy between the reference image and warped image. JDACS [26] proposes cross-view consistency of extracted semantic features and provides supervision of the segmentation map by non-negative matrix factorization. However, it still requires a pretrained semantic feature extraction backbone, and suffers from unstable convergence of cross-view semantic consistency that fails to provide reliable supervision. U-MVSNet [27] employs a flow-depth consistency loss to resolve ambiguous supervision. The dense 2D optical flow correspondences are used to generate pseudo labels for uncertainty-aware consistency, which improves the supervision to some extent. However, this method cannot be trained in an end-toend fashion, as it requires complex pre-training and fine-tuning. Self-supervised CVP-MVSNet [31] also proposes to learn depth pseudo labels with unsupervised pre-training, followed with the iterative self-supervised training to refine the pseudo label; however, it is still affected by ambiguous supervision. Furthermore, these unsupervised and self-supervised methods lack an occlusion-aware module to learn features from different viewing directions, leading to incomplete reconstruction of point clouds.

In contrast, our simple but effective model directly learns the geometric features of the scene by a reference view synthesis loss. This significantly reduces the training complexity and alleviates ambiguous photometric supervision. It also avoids the problems caused by occlusions through NeRF-like rendering.

2.3 Multi-View Neural Rendering

Recently, various neural rendering methods [3,13,15,20,38] have been presented, focusing on task of novel view synthesis. In particular, Neural Radiance Fields [15] represent scenes with a continuous implicit function of positions and orientations



Fig. 3. Overview of our unsupervised MVS approach (RC-MVSNet). a) Unsupervised Backbone CasMVSNet predicts initial depth map by photometric consistency and provides depth priors for rendering consistency network. b) Rendering Consistency Network generates image and depth by neural rendering under the guidance of depth priors. c) The rendered image is supervised by the reference view synthesis loss. d) The rendered depth is supervised by the depth rendering consistency loss.

for high quality view synthesis; several following up works [16,35,36,37] improve both its efficiency and performance. There are also a few extensions [4,18,24] of NeRF which introduce multi-view information to enhance the generalization ability of view synthesis. MVSNeRF [4] proposes to leverage plane-swept cost volumes, which have been widely used in multi-view stereo, for geometry-aware scene understanding, and combines this with volumetric rendering for neural radiance field reconstruction; however, it fails to generate high-quality depth prediction in an unsupervised manner. NerfingMVS [24] uses sparse depth points from SfM for learning depth completion to guide the optimization process of NeRF. Our method takes advantage of both precise neural rendering and strong generalization of cost volumes and provides accurate depth estimation based on end-to-end unsupervised learning, which surpasses all previous unsupervised multi-view-stereo approaches and demonstrates accurate reconstruction of both indoor objects and out-of-distribution outdoor scenes.

3 Method

In this section, we describe RC-MVSNet, our proposed method for unsupervised multi-view stereo. Given N images as input, with their corresponding cameras' intrinsic and extrinsic parameters, our method predicts a depth map in the reference camera view. The overall pipeline is illustrated in Fig. 3. It consists of a backbone branch and an auxiliary branch. The backbone is built upon CasMVS-Net [8], which predicts a depth map in a coarse-to-fine manner, as described in

Sec. 3.1. The auxiliary branch is built upon the neural radiance fields, which aside from geometry also models view-dependent scene appearance. Using volumetric rendering, the reference view is synthesized and compared to the input image (in Sec. 3.2), resulting in accurate supervision even in the case of view-dependent effects. To ensure the geometric consistency between the network branches, an additional depth rendering consistency loss is introduced in Sec. 3.3, leading to more complete and accurate depth predictions, even in the case of occlusions. Note that during training, these two branches are simultaneously optimized, providing supervision to each other. During inference, only the backbone is used to obtain the depth prediction. We present the network optimization in detail in Sec. 3.4.

3.1 Unsupervised Multi-view Stereo Network

The backbone network closely follows the CasMVSNet [8] architecture. Input images $\{I_j\}_{j=1}^N$ are initially encoded with a shared 2D U-Net that generates pixelwise features. Afterwards, a feature cost volume is constructed in the reference camera frustum. Each volume voxel's position is projected into every input image using cameras' intrinsic and extrinsic parameters, and pixel features are queried via bilinear interpolation. This results in warped feature volumes $\{V_j\}_{j=1}^N$ that are fused across views into a common feature volume C by computing feature variance:

$$C = Var\left(V_1, \cdots, V_N\right) = \frac{\sum_{j=1}^{N} \left(V_j - \overline{V}_j\right)^2}{N}.$$
(1)

The feature cost volume C is further refined with a 3D U-Net and finally outputs a depth map in the reference view in a coarse-to-fine manner. Further details can be found in the supplementary material.

Photometric Consistency. To supervise depth map prediction without any ground truth depth, existing methods [6,9,11,27,31] enforce photometric consistency between reference and other source input views. The key idea is to maximize the similarity between the reference image I_1 and any source image I_j after being warped to the reference view. Given corresponding intrinsics K, the relative transformation from reference to source view T, and predicted depth map D in the reference view, we warp any pixel p_i in the reference image (in a homogeneous notation) to the source image via inverse warping:

$$\hat{p}_i = KT \left(D(p_i) \cdot K^{-1} p_i \right). \tag{2}$$

The warped source image \hat{I}_1^j is synthesized by bilinearly sampling the source image's color values at the warped pixel location \hat{p}_i , i.e. $\hat{I}_1^j(p_i) = I_j(\hat{p}_i)$. In additon to the warped image \hat{I}_1^j , a binary mask M_j is also generated, masking out invalid pixels that are projected outside the source image bounds. In the unsupervised MVS system, all N-1 source views are warped into the reference view to compute the photometric consistency loss:

$$\mathcal{L}_{PC} = \sum_{j=2}^{N} \frac{1}{\|M_{j}\|_{1}} \left(\left\| \left(\hat{I}_{1}^{j} - I_{1} \right) \odot M_{j} \right\|_{2} + \left\| \left(\nabla \hat{I}_{1}^{j} - \nabla I_{1} \right) \odot M_{j} \right\|_{2} \right).$$
(3)

Here, ∇ refers to a pixel-wise gradient and \odot represents pixel-wise product.

3.2 Reference View Synthesis

Photometric consistency assumes consistent pixel colors across different input views, which does not hold in the case of view-dependent effects that are common with reflective materials. Furthermore, if certain pixels in the reference view are occluded in the matched source view, incorrect color similarity is enforced. To cope with these shortcomings, instead of using error-prone image warping of source images, we propose to learn the synthesis of reference views from source images using neural radiance fields [15], with the network learning how to take care of view-dependent effects and occlusions.

Implicit Neural Volume Construction. Similar to the cost volume C for the backbone branch, an implicit neural volume C' is also constructed via computation of feature variance, as defined in Eq. (1). However, since we want to synthesize the reference view image using *only* the information from the source views, we calculate the variance volume C' from warped volumes $\{V_j\}_{j=2}^N$, without the reference feature volume V_1 . We thus gather features across N-1 source views, and apply a 3D U-Net $U(\cdot)$ to learn an implicit neural volume F that we denote as F = U(C').

Point Feature Aggregation. To restore the scene geometry and appearance, rays are emitted from the reference camera center in the viewing direction of ϕ . For each point $q \in \mathbb{R}^3$ along the ray we query RGB values $\ell = \{I_j[x_j, y_j]\}_{j=2}^N$ by bilinearly sampling source images, where x_j and y_j represent the 2D location of point q, when it is projected to the image coordinate system of j^{th} view. In addition to these properties, we also obtain a neural feature f = F(q) from the implicit neural volume F by trilinear interpolation at the location q. These neural features overcome the poor generalization of traditional neural radiance field approaches that require per-scene training or fine-tuning, by regressing properties from the geometric-aware cost volume (e.g. implicit neural volume in our case). This improves our method's robustness on datasets that are very different from the training set. These input features are concatenated into point features and we use an MLP $M(\cdot)$ to convert these features into predicted volume density σ and color c = [r, g, b]. The whole process can be summarized as:

$$\sigma, c = M(q, \phi, f, \ell). \tag{4}$$

Following previous works [4,15,17], we encode the location q and direction d using positional encoding.

Reference View Synthesis Loss. As in NeRF [15], we sample K point candidates along the camera ray $\mathbf{r}(t) = \mathbf{o} + t\phi$, where \mathbf{o} origin is the camera and ϕ is the pixel's viewing direction. For the sample distance t_k in the range from near to far plane ($t_n \leq t_k \leq t_f$), we query the corresponding density σ_k and color c_k . To compute pixel-wise color $\hat{\mathbf{C}}(\mathbf{r})$, color values are integrated along the ray $\mathbf{r}(t)$:

$$\hat{\mathbf{C}}(\mathbf{r}) = \sum_{k=1}^{K} w_k c_k.$$
(5)

Here $w_k = T_k (1 - \exp(-\sigma_k \delta_k))$ is the integration weight, where $\delta_k = t_{k+1} - t_k$ refers to the interval between adjacent sampling locations and T_k indicates the accumulation of transmittance from t_n to t_f :

$$T_k = \exp\left(-\sum_{k'=1}^k \sigma_{k'} \delta_{k'}\right). \tag{6}$$

To optimize the radiance field color output $\hat{\mathbf{C}}(\mathbf{r})$ in Eq. (5), i.e. the synthesized RGB value of a pixel (x, y) in the reference view, the rendering consistency loss L_{RC} composed of a mean squared error is applied:

$$\mathcal{L}_{RC} = \|\hat{\mathbf{C}}(\mathbf{r}) - \mathbf{C}(\mathbf{r})\|_2^2.$$
(7)

Here $\mathbf{C}(\mathbf{r})$ refers to pixel color in the reference image $I_1[x, y]$.

3.3 Depth Rendering Consistency

Using an auxiliary rendering branch and applying rendering consistency loss significantly improves the quality of unsupervised depth predictions, as can be seen in our ablations in Sec. 4.4. This results from joint training of both branches and optimizing a shared 2D image encoder. To further benefit from rendering consistency and additionally improve unsupervised depth prediction, we also apply Gaussian-Uniform mixture sampling and a depth rendering consistency loss, which increase the beneficial coupling of both network branches.

Gaussian-Uniform Mixture Sampling. To sample points along the ray for volumetric rendering in Eq. (5), existing methods [4,15,18,24,36] partition $[t_n, t_f]$ into K bins following a uniform distribution:

$$t_k \sim \mathcal{U}\left[t_n + \frac{k-1}{K}\left(t_f - t_n\right), t_n + \frac{k}{K}\left(t_f - t_n\right)\right].$$
(8)

This sampling strategy is inefficient and also independent of depth predictions of the backbone branch. We propose to sample the point candidates following a Gaussian distribution with the guidance of the depth prior from predicted depth map D, as illustrated in Fig. 4. Assuming the predicted depth for a pixel

 $\mathbf{p}=(x,y)$ to be $z_{\mathbf{p}}=D(x,y),$ we sample the candidates using the following distribution:

$$t_k \sim \mathcal{N}\left(z_{\mathbf{p}}, s_{\mathbf{p}}^2\right). \tag{9}$$

where
$$s_{\mathbf{p}} = \frac{\min(|z_{\mathbf{p}} - t_f|, |z_{\mathbf{p}} - t_n|)}{3}.$$
 (10)

Here $z_{\mathbf{p}}$ and $s_{\mathbf{p}}$ are the mean and standard deviation of the proposed normal distribution, respectively. This enables geometry features for rendering to be more efficiently optimized, since we sample more point candidates close to the object surface. Furthermore, since we employ differentiable Gaussian sampling, the depth predictions from the backbone branch obtain useful cues on how to improve depth, also in the case of occlusions. In order to make the training stable and help the network to converge, particularly during the beginning stage of the end-to-end training, half of the samples are drawn from the Gaussian distribution from Eq. (9) and the other half is uniformly distributed between the near and far planes.



Fig. 4. Demonstration of our proposed Gaussian-Uniform mixture sampling.

Depth Rendering Consistency Loss. Our rendering consistency network encodes geometry through density estimation. We can convert the density estimate to a depth map in reference view through density integration along the ray:

$$\hat{z}(\mathbf{r}) = \sum_{k=1}^{K} w_k t_k.$$
(11)

We use the same rendering weights w_k as in Eq. (5), and t_k is a sampled point candidate as introduced in Sec. 3.2, following the normal distribution. To enforce that geometry estimated by the backbone branch and by the auxiliary rendering branch match, we minimize the difference between the depth maps with an additional depth rendering consistency loss L_{DC} , defined as a Smooth-L1 loss [7]:

$$\mathcal{L}_{DC} = \mathrm{smooth}_{L_1} \left(\hat{z}(\mathbf{r}) - z_{\mathbf{p}} \right). \tag{12}$$

Here, $\hat{z}(\mathbf{r})$ is the integrated depth value from the rendering consistency network and $z_{\mathbf{p}}$ the corresponding pixel depth value from the backbone network.

3.4 End-to-end Network Optimization

Unlike pseudo-label-based multi-stage self-supervised methods [27,31], our model is trained from scratch in an end-to-end fashion, without any pre-processing or pre-training. The two parts indicated in Sec. 3, the backbone and auxiliary network branches, are jointly trained on the whole dataset by minimizing the following overall loss:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{PC} + \lambda_2 \mathcal{L}_{RC} + \lambda_3 \mathcal{L}_{DC} + \lambda_4 \mathcal{L}_{SSIM} + \lambda_5 \mathcal{L}_{Smooth} + \lambda_6 \mathcal{L}_{DA}.$$
 (13)

Losses \mathcal{L}_{SSIM} and \mathcal{L}_{Smooth} , commonly used in previous works [6,9,11,26], ensure structural similarity and smoothness of predicted depth maps. As proposed by [26], we also apply color fluctuation augmentation to input images and calculate the L1 loss between depth maps predicted under different color augmentations, resulting in L_{DA} . We use fixed weights $\lambda_1 = 0.8$, $\lambda_2 = 1.0$, $\lambda_3 = 1.0$, $\lambda_4 = 0.2$ and $\lambda_5 = 0.0067$. The weight for data augmentation is initialized with $\lambda_6 = 0.01$ and doubled every two epochs. We evaluate the effect of different losses on depth prediction performance in the ablation study (see Sec. 4.4).

4 Experiments

4.1 Datasets

The **DTU** dataset [2] is an indoor dataset with multi-view images and corresponding camera parameters. We follow the settings of MVSNet [33] for dividing training, validation and evaluation sets. There are 124 scenes that are scanned from 49 or 64 views under 7 different light conditions. In the DTU benchmark, models are trained on training set and tested on the evaluation set. We use the official error metrics in DTU to evaluate the Accuracy, which is measured as the distance from the result to the structured light reference, indicating the quality of the reconstruction, and *Completeness*, which is measured as the distance from the ground truth reference to the reconstructed result, indicating the proportion of the reconstructed part in the entire point cloud. Overall is the average of Ac*curacy* and *Completeness* and it reflects the overall quality of the reconstruction. Tanks and Temples [12] is a large-scale dataset with various outdoor scenes. It contains an intermediate subset and an advanced subset. The evaluation on this benchmark is conducted online by submitting generated point clouds to the official website. In this benchmark, F-score is calculated for each scene and we compare the mean F-score of intermediate and advanced subset respectively.

4.2 Implementation Details

Training Details. The proposed RC-MVSNet is trained on the DTU dataset. Following [6,9,11,26], we use the high-resolution DTU data provided by the open source code of MVSNet [33]. We first resize the input images to 600×800 ,

Table 1. Point cloud evaluation results on DTU [2]. The lower is better for Accuracy (Acc.), Completeness (Comp.), and Overall. The sections are partitioned into supervised, multi-stage self-supervised and end-to-end (E2E) unsupervised, respectively. The best result is highlighted in bold for each category. All the results other than ours are from previously published literature.

$\begin{tabular}{ c c c c c c c c c c c c c c c c c c c$		Method	$\mathrm{Acc.}{\downarrow}$	Comp.↓	$Overall.\downarrow$
$ \begin{array}{c} \mbox{MVSNet [33]} & 0.396 & 0.527 & 0.462 \\ \mbox{Cas-MVSNet [8]} & 0.325 & 0.385 & 0.355 \\ \mbox{Supervised} & \mbox{PatchmatchNet [21]} & 0.427 & 0.277 & 0.352 \\ \mbox{CVP-MVSNet [32]} & {\bf 0.296} & 0.406 & 0.351 \\ \mbox{UCSNet [5]} & 0.338 & 0.349 & 0.344 \\ \mbox{GBi-Net [14]} & 0.315 & {\bf 0.262} & {\bf 0.289} \\ \mbox{Multi-Stage Self_sup CVP-MVSNet [31]} & {\bf 0.308} & 0.418 & 0.363 \\ \mbox{Self-Sup.} & \mbox{U-MVSNet [27]} & 0.354 & {\bf 0.3535} & {\bf 0.3535} \\ \mbox{Unsup_MVSNet [11]} & 0.881 & 1.073 & 0.977 \\ \mbox{MVS2 [6]} & 0.760 & 0.515 & 0.637 \\ \mbox{E2E UnSup.} & \mbox{M3VSNet [9]} & 0.636 & 0.531 & 0.583 \\ \mbox{JDACS-MS [26]} & 0.398 & 0.318 & 0.358 \\ \mbox{RC-MVSNet} & {\bf 0.396} & 0.295 & 0.345 \\ \end{tabular}$		SurfaceNet [10]	0.450	1.040	0.745
Cas-MVSNet [8] 0.325 0.385 0.355 Supervised PatchmatchNet [21] 0.427 0.277 0.352 CVP-MVSNet [32] 0.296 0.406 0.351 UCSNet [5] 0.338 0.349 0.344 GBi-Net [14] 0.315 0.262 0.289 Multi-Stage Self-sup CVP-MVSNet [31] 0.308 0.418 0.363 Self-Sup. U-MVSNet [27] 0.354 0.3535 0.3535 MVS2 [6] 0.760 0.515 0.637 E2E Unsup_MVSNet [9] 0.636 0.531 0.583 JDACS-MS [26] 0.398 0.318 0.358 RC-MVSNet 0.396 0.295 0.345		MVSNet [33]	0.396	0.527	0.462
Supervised PatchmatchNet [21] 0.427 0.277 0.352 CVP-MVSNet [32] 0.296 0.406 0.351 UCSNet [5] 0.338 0.349 0.344 GBi-Net [14] 0.315 0.262 0.289 Multi-Stage Self_sup CVP-MVSNet [31] 0.308 0.418 0.363 Self-Sup. U-MVSNet [27] 0.354 0.3535 0.3535 VInsup_MVSNet [27] 0.354 0.3535 0.3535 E2E UnSup. M3VSNet [9] 0.636 0.531 0.583 JDACS-MS [26] 0.398 0.318 0.358 RC-MVSNet 0.396 0.295 0.345		Cas-MVSNet [8]	0.325	0.385	0.355
CVP-MVSNet [32] 0.296 0.406 0.351 UCSNet [5] 0.338 0.349 0.344 GBi-Net [14] 0.315 0.262 0.289 Multi-Stage Self_sup CVP-MVSNet [31] 0.308 0.418 0.363 Self-Sup. U-MVSNet [27] 0.354 0.3535 0.3535 Versup_MVSNet [11] 0.881 1.073 0.977 MVS2 [6] 0.760 0.515 0.637 E2E UnSup. M3VSNet [9] 0.636 0.531 0.583 JDACS-MS [26] 0.398 0.318 0.358 RC-MVSNet 0.396 0.295 0.345	Supervised	PatchmatchNet [21]	0.427	0.277	0.352
UCSNet [5] 0.338 0.349 0.344 GBi-Net [14] 0.315 0.262 0.289 Multi-Stage Self_sup CVP-MVSNet [31] 0.308 0.418 0.363 Self-Sup. U-MVSNet [27] 0.354 0.3535 0.3535 Winsup_MVSNet [27] 0.354 0.3535 0.3535 E2E UnSup. M3VSNet [9] 0.636 0.531 0.583 JDACS-MS [26] 0.398 0.318 0.358 RC-MVSNet 0.396 0.295 0.345		CVP-MVSNet [32]	0.296	0.406	0.351
GBi-Net [14] 0.315 0.262 0.289 Multi-Stage Self_sup CVP-MVSNet [31] 0.308 0.418 0.363 Self-Sup. U-MVSNet [27] 0.354 0.3535 0.3535 Unsup_MVSNet [27] 0.354 0.3535 0.3535 0.3535 E2E Unsup_MVSNet [11] 0.881 1.073 0.977 MVS2 [6] 0.760 0.515 0.637 E2E Unsup_MVSNet [9] 0.636 0.531 0.583 JDACS-MS [26] 0.398 0.318 0.3585 RC-MVSNet 0.396 0.295 0.345		UCSNet [5]	0.338	0.349	0.344
Multi-Stage Self_sup CVP-MVSNet [31] 0.308 0.418 0.363 Self-Sup. U-MVSNet [27] 0.354 0.3535 0.3535 0.3535 Winsup_MVSNet [11] 0.881 1.073 0.977 MVS2 [6] 0.760 0.515 0.637 E2E Unsup_MVSNet [9] 0.636 0.531 0.583 JDACS-MS [26] 0.398 0.318 0.358 RC-MVSNet 0.396 0.295 0.345		GBi-Net [14]	0.315	0.262	0.289
Self-Sup. U-MVSNet [27] 0.354 0.3535 0.3535 Unsup_MVSNet [11] 0.881 1.073 0.977 MVS2 [6] 0.760 0.515 0.637 E2E UnSup. M3VSNet [9] 0.636 0.531 0.583 JDACS-MS [26] 0.398 0.318 0.358 RC-MVSNet 0.396 0.295 0.345	Multi-Stage	Self_sup CVP-MVSNet [31]	0.308	0.418	0.363
Unsup_MVSNet [11] 0.881 1.073 0.977 MVS2 [6] 0.760 0.515 0.637 E2E Unsup. M3VSNet [9] 0.636 0.531 0.583 JDACS-MS [26] 0.398 0.318 0.358 RC-MVSNet 0.396 0.295 0.345	Self-Sup.	U-MVSNet [27]	0.354	0.3535	0.3537
MVS2 [6] 0.760 0.515 0.637 E2E UnSup. M3VSNet [9] 0.636 0.531 0.583 JDACS-MS [26] 0.398 0.318 0.358 RC-MVSNet 0.396 0.295 0.345		Unsup_MVSNet [11]	0.881	1.073	0.977
E2E UnSup. M3VSNet [9] 0.636 0.531 0.583 JDACS-MS [26] 0.398 0.318 0.358 RC-MVSNet 0.396 0.295 0.345		MVS2 [6]	0.760	0.515	0.637
JDACS-MS [26] 0.398 0.318 0.358 RC-MVSNet 0.396 0.295 0.345	E2E UnSup.	M3VSNet [9]	0.636	0.531	0.583
RC-MVSNet 0.396 0.295 0.345		JDACS-MS [26]	0.398	0.318	0.358
		RC-MVSNet	0.396	0.295	0.345

following previous methods. Then we crop resized images into 512×640 patches. We adopt the backbone of Cas-MVSNet [8] to construct our multi-scale pipeline with 3 stages. For each stage, we use different feature maps and the 3D-CNN network parameters. The whole network is optimized by an Adam optimizer in Pytorch for 15 epochs with an initial learning rate of 0.0001, which is down-scaled by a factor of 2 after 10, 12, and 14 epochs. We train with a batch size of 4 using four NVIDIA RTX 3090 GPUs.

Testing Details. The model trained on DTU training set is used for testing on DTU testing set. The input image number N is set to 5, each with a resolution of 1152×1600 . It takes 0.475 seconds to test each sample. The model trained on DTU training dataset is directly used for testing on Tanks and Temples intermediate and advanced datasets without finetuning. The image sizes are set to 1024×1920 or 1024×2048 and the input image number N is set to 7. We then filter predicted depth maps of a scene using photometric and geometric consistencies and fuse then into one point cloud. We define geometric and photometric consistencies similarly to the ones used in MVSNet [33], a detailed description can be found in supplemental material. Fig. 1 shows visualizations of point clouds of RC-MVSNet and previous unsupervised methods, and Fig. 5 shows reconstructions of our method on the test dataset.

4.3 Benchmark Performance

Evaluation on DTU Dataset. We evaluate the depth prediction performance on the DTU test set, and compare with previous state-of-the-art methods in Table 1. Our RC-MVSNet architecture achieves the best accuracy, completeness and overall score (lower is better for all metrics) among all end-to-end

Table 2. Point cloud evaluation results on the intermediate subsets of Tanks and Temples dataset [12]. Higher scores are better. The Mean is the average score of all scenes. The sections are partitioned into supervised, multi-stage self-supervised and end-to-end (E2E) unsupervised, respectively. The best result is highlighted in bold for every section.

		Tanks&Temples intermediate								
	Method	$\operatorname{Mean}\uparrow$	Family↑	Francis↑	Horse↑	Lighthouse↑	$M60\uparrow$	$Panther\uparrow$	Playground↑	Train↑
	MVSNet [33]	43.48	55.99	28.55	25.07	50.79	53.96	50.86	47.90	34.69
	CIDER [28]	46.76	56.79	32.39	29.89	54.67	53.46	53.51	50.48	42.85
Supervised	PatchmatchNet [21]	53.15	66.99	52.64	43.24	54.87	52.87	49.54	54.21	50.81
	CVP-MVSNet [32]	54.03	76.50	47.74	36.34	55.12	57.28	54.28	57.43	47.54
	UCSNet [5]	54.83	76.09	53.16	43.03	54.00	55.60	51.49	57.38	47.89
	Cas-MVSNet [8]	56.42	76.36	58.45	46.20	55.53	56.11	54.02	58.17	46.56
Multi-Stage	Self_sup CVP-MVSNet [31]	46.71	64.95	38.79	24.98	49.73	52.57	51.53	50.66	40.45
Self-Sup.	U-MVSNet [27]	57.15	76.49	60.04	49.20	55.52	55.33	51.22	56.77	52.63
	MVS2 [6]	37.21	47.74	21.55	19.50	44.54	44.86	46.32	43.38	29.72
E2E Unsup.	M3VSNet [9]	37.67	47.74	24.38	18.74	44.42	43.45	44.95	47.39	30.31
	JDACS-MS [26]	45.48	66.62	38.25	36.11	46.12	46.66	45.25	47.69	37.16
	RC-MVSNet	55.04	75.26	53.50	45.52	53.49	54.85	52.30	56.06	49.37

unsupervised methods. Our model improves the overall score from 0.358 of JDACS [26] to 0.345. The completeness is improved by 0.023, demonstrating the effectiveness of our method to resolve incomplete predictions caused by occlusions. The overall score is also better than the multi-stage self-supervised approaches, and even most supervised methods. In addition to the performance of point cloud reconstruction, we also provide evaluations of predicted depth maps. We compare them with supervised MVSNet [33] and previous unsupervised approaches [6,9,11,26], the results are summarized in Table. 3. Our method again achieves state-of-the-art accuracy. Visualizations of our depth predictions can be found in the supplementary material.



Fig. 5. Point cloud visualization of our method on DTU [2] and Tanks and Temples [12].

Evaluation on Tanks and Temples. We train the proposed RC-MVSNet on the DTU train set, and test on the Tanks and Temples dataset without finetuning. We compare our method to state-of-the-art supervised, pseudo-label-based multi-stage self-supervised methods and end-to-end unsupervised methods. Table 2 shows the evaluation results on the intermediate subset. Our RC-MVSNet achieves the best performance among unsupervised methods. In particular, our

Method	$<\!2\uparrow$	$<4\uparrow$	$<\!8\uparrow$
MVSNet [33]	0.704	0.778	0.815
MVSNeRF [21]	0.510	0.645	0.734
Unsup_MVSNet [11]	0.317	0.384	0.402
M3VSNet [9]	0.603	0.769	0.857
JDACS-MS [26]	0.553	0.705	0.786
RC-MVSNet	0.730	0.795	0.863

Table 3. Depth map evaluation results in terms of accuracy on DTU evaluation set [2] (higher is better). All thresholds are given in millimeters.

mean score achieves significant improvement of +17.37 over M3VSNet [9] and +9.56 over JDACS [26]. Moreover, we also obtain state-of-the-art performance in all sub-scenes, which fully confirms the effectiveness of our method. Our **anonymous** evaluation on the leaderboard [1] is named as RC-MVSNet. The results of advanced subset can be found in the supplementary material, along with additional qualitative results and discussion.

4.4 Ablation Study

We provide an ablation analysis, demonstrating the point cloud reconstruction performance gain of each component of our proposed method. We also evaluate performance when sampling different number of rays during rendering in the supplementary material.



Fig. 6. Qualitative results of RC-MVSNet on scan33 (top) and scan48 (bottom) of the DTU dataset.

Effect of each component of the unsupervised approach. To evaluate the performance gain of each part in our method, we provide qualitative comparison of models trained using different combinations of core contributions in Fig. 6. We can easily observer that both reference view synthesis and depth rendering consistency yield better reconstruction results, especially when considering completeness. Quantitative results are listed in Table. 4, further confirming the importance of all proposed network components.

Table 4. Ablation study of different components of our proposed unsupervised approach (the lower is better). L_{PC} : Photometric Consistency Loss. L_{DA} : Data Augmentation Loss. L_{RC} : Reference View Synthesis Loss. G-U: Gaussian-Uniform Mixture Sampling. L_{DC} : Depth Rendering Consistency Loss. The last three are contributions in this work.

L_{PC}	L_{DA}	L_{RC}	G-U	L_{DC}	$Acc\downarrow$	$\mathrm{Comp}{\downarrow}$	Overall↓
\checkmark					0.462	0.328	0.395
\checkmark	\checkmark				0.419	0.346	0.383
\checkmark	\checkmark	\checkmark			0.415	0.300	0.357
\checkmark	\checkmark	\checkmark	\checkmark		0.404	0.299	0.352
\checkmark	\checkmark	\checkmark		\checkmark	0.397	0.299	0.348
\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	0.396	0.295	0.345

5 Conclusions

We show that unsupervised multi-view-stereo can be improved by challenging the previous assumption of photometric consistency. We thus propose a novel unsupervised MVS approach based on rendering consistency (RC-MVSNet). To handle the ambiguous supervision, we propose a reference view synthesis loss via differentiable volumetric rendering. To solve the incompleteness caused by occlusions, we introduce Gaussian-Uniform mixture sampling to learn geometry features close to the object surface. To further improve the robustness and smoothness of the depth map, we propose a depth rendering consistency loss. The experiments demonstrate the effectiveness of our RC-MVSNet approach.

6 Acknowledgments

This project is funded by the ERC Starting Grant Scan2CAD (804724), a TUM-IAS Rudolf Mößbauer Fellowship, and the German Research Foundation (DFG) Grant Making Machine Learning on Static and Dynamic 3D Data Practical.

References

- 1. Tanks and temples leaderboard. https://www.tanksandtemples.org/ leaderboard 13
- Aanæs, H., Jensen, R.R., Vogiatzis, G., Tola, E., Dahl, A.B.: Large-scale data for multiple-view stereopsis. International Journal of Computer Vision 120(2), 153– 168 (2016) 1, 3, 10, 11, 12, 13
- Bi, S., Xu, Z., Sunkavalli, K., Hašan, M., Hold-Geoffroy, Y., Kriegman, D., Ramamoorthi, R.: Deep reflectance volumes: Relightable reconstructions from multiview photometric images. In: European Conference on Computer Vision. pp. 294– 311. Springer (2020) 4
- Chen, A., Xu, Z., Zhao, F., Zhang, X., Xiang, F., Yu, J., Su, H.: Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. arXiv preprint arXiv:2103.15595 (2021) 3, 5, 7, 8
- Cheng, S., Xu, Z., Zhu, S., Li, Z., Li, L.E., Ramamoorthi, R., Su, H.: Deep stereo using adaptive thin volume representation with uncertainty awareness. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2524–2534 (2020) 11, 12
- 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) 2, 6, 10, 11, 12
- Girshick, R.: Fast r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 1440–1448 (2015) 9
- 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: CVPR (2020) 1, 4, 5, 6, 11, 12
- Huang, B., Yi, H., Huang, C., He, Y., Liu, J., Liu, X.: M3vsnet: Unsupervised multi-metric multi-view stereo network. In: 2021 IEEE International Conference on Image Processing (ICIP). pp. 3163–3167. IEEE (2021) 2, 6, 10, 11, 12, 13
- Ji, M., Gall, J., Zheng, H., Liu, Y., Fang, L.: Surfacenet: An end-to-end 3d neural network for multiview stereopsis. In: ICCV (2017) 11
- 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) 2, 4, 6, 10, 11, 12, 13
- Knapitsch, A., Park, J., Zhou, Q.Y., Koltun, V.: Tanks and temples: Benchmarking large-scale scene reconstruction. ACM ToG 36(4), 1–13 (2017) 1, 3, 10, 12
- Lombardi, S., Simon, T., Saragih, J., Schwartz, G., Lehrmann, A., Sheikh, Y.: Neural volumes: Learning dynamic renderable volumes from images. arXiv preprint arXiv:1906.07751 (2019) 4
- 14. Mi, Z., Chang, D., Xu, D.: Generalized binary search network for highly-efficient multi-view stereo (2021) 1, 11
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: European conference on computer vision. pp. 405–421. Springer (2020) 3, 4, 7, 8
- Rebain, D., Jiang, W., Yazdani, S., Li, K., Yi, K.M., Tagliasacchi, A.: Derf: Decomposed radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14153–14161 (2021) 5
- Roessle, B., Barron, J.T., Mildenhall, B., Srinivasan, P.P., Nießner, M.: Dense depth priors for neural radiance fields from sparse input views. arXiv preprint arXiv:2112.03288 (2021) 7

- 16 D. Chang et al.
- 18. Rosu, R.A., Behnke, S.: Neuralmvs: Bridging multi-view stereo and novel view synthesis. arXiv preprint arXiv:2108.03880 (2021) 3, 5, 8
- Seitz, S.M., Curless, B., Diebel, J., Scharstein, D., Szeliski, R.: A comparison and evaluation of multi-view stereo reconstruction algorithms. In: 2006 IEEE computer society conference on computer vision and pattern recognition (CVPR'06). vol. 1, pp. 519–528. IEEE (2006) 1
- Thies, J., Zollhöfer, M., Nießner, M.: Deferred neural rendering: Image synthesis using neural textures. ACM Transactions on Graphics (TOG) 38(4), 1–12 (2019)
 4
- Wang, F., Galliani, S., Vogel, C., Speciale, P., Pollefeys, M.: Patchmatchnet: Learned multi-view patchmatch stereo. In: CVPR (2021) 1, 4, 11, 12, 13
- Wang, Q., Wang, Z., Genova, K., Srinivasan, P.P., Zhou, H., Barron, J.T., Martin-Brualla, R., Snavely, N., Funkhouser, T.: Ibrnet: Learning multi-view image-based rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4690–4699 (2021) 3
- 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) 4
- Wei, Y., Liu, S., Rao, Y., Zhao, W., Lu, J., Zhou, J.: Nerfingmvs: Guided optimization of neural radiance fields for indoor multi-view stereo. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5610–5619 (2021) 5, 8
- 25. Wei, Z., Zhu, Q., Min, C., Chen, Y., Wang, G.: Aa-rmvsnet: Adaptive aggregation recurrent multi-view stereo network. In: ICCV (October 2021) 4
- 26. Xu, H., Zhou, Z., Qiao, Y., Kang, W., Wu, Q.: Self-supervised multi-view stereo via effective co-segmentation and data-augmentation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 2, p. 6 (2021) 2, 3, 4, 10, 11, 12, 13
- Xu, H., Zhou, Z., Wang, Y., Kang, W., Sun, B., Li, H., Qiao, Y.: Digging into uncertainty in self-supervised multi-view stereo. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6078–6087 (2021) 4, 6, 10, 11, 12
- Xu, Q., Tao, W.: Learning inverse depth regression for multi-view stereo with correlation cost volume. In: AAAI. vol. 34 (2020) 12
- 29. Xu, Q., Tao, W.: Pvsnet: Pixelwise visibility-aware multi-view stereo network (2020) 1
- 30. Yan, J., Wei, Z., Yi, H., Ding, M., Zhang, R., Chen, Y., Wang, G., Tai, Y.W.: Dense hybrid recurrent multi-view stereo net with dynamic consistency checking. In: ECCV (2020) 4
- Yang, J., Alvarez, J.M., Liu, M.: Self-supervised learning of depth inference for multi-view stereo. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7526–7534 (2021) 4, 6, 10, 11, 12
- Yang, J., Mao, W., Alvarez, J.M., Liu, M.: Cost volume pyramid based depth inference for multi-view stereo. In: CVPR (2020) 4, 11, 12
- 33. Yao, Y., Luo, Z., Li, S., Fang, T., Quan, L.: Mvsnet: Depth inference for unstructured multi-view stereo. In: ECCV (2018) 1, 4, 10, 11, 12, 13
- 34. Yao, Y., Luo, Z., Li, S., Shen, T., Fang, T., Quan, L.: Recurrent mysnet for high-resolution multi-view stereo depth inference. In: CVPR (2019) 4
- Yu, A., Fridovich-Keil, S., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. arXiv preprint arXiv:2112.05131 (2021) 5

- 36. Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelnerf: Neural radiance fields from one or few images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4578–4587 (2021) 5, 8
- Zhang, K., Riegler, G., Snavely, N., Koltun, V.: Nerf++: Analyzing and improving neural radiance fields. arXiv preprint arXiv:2010.07492 (2020) 5
- Zhou, T., Tucker, R., Flynn, J., Fyffe, G., Snavely, N.: Stereo magnification: Learning view synthesis using multiplane images. arXiv preprint arXiv:1805.09817 (2018)
 4