

Enhanced Stable View Synthesis

Nishant Jain*
Indian Institute of Technology
Roorkee, India
njain@cs.iitr.ac.in

Suryansh Kumar*[†]
ETH Zürich
Switzerland
sukumar@ethz.ch

Luc Van Gool
ETH Zürich
Switzerland
vangool@ethz.ch

Abstract

We introduce an approach to enhance the novel view synthesis from images taken from a freely moving camera. The introduced approach focuses on outdoor scenes where recovering accurate geometric scaffold and camera pose is challenging, leading to inferior results using the state-of-the-art stable view synthesis (SVS) method. SVS and related methods fail for outdoor scenes primarily due to (i) over-relying on the multiview stereo (MVS) for geometric scaffold recovery and (ii) assuming COLMAP computed camera poses as the best possible estimates, despite it being well-studied that MVS 3D reconstruction accuracy is limited to scene disparity and camera-pose accuracy is sensitive to key-point correspondence selection. This work proposes a principled way to enhance novel view synthesis solutions drawing inspiration from the basics of multiple view geometry. By leveraging the complementary behavior of MVS and monocular depth, we arrive at a better scene depth per view for nearby and far points, respectively. Moreover, our approach jointly refines camera poses with image-based rendering via multiple rotation averaging graph optimization. The recovered scene depth and the camera-pose help better view-dependent on-surface feature aggregation of the entire scene. Extensive evaluation of our approach on the popular benchmark dataset, such as Tanks and Temples, shows substantial improvement in view synthesis results compared to the prior art. For instance, our method shows **1.5 dB** of PSNR improvement on the Tank and Temples. Similar statistics are observed when tested on other benchmark datasets such as FVS, Mip-NeRF 360, and DTU.

1. Introduction

Image-based rendering, popularly re-branded as view synthesis, is a long-standing problem in computer vision and graphics [42, 44]. This problem aims to develop a

method that allows the user to seamlessly explore the scene via rendering of the scene from a sparse set of captured images [2, 20, 42]. Furthermore, the rendered images must be as realistic as possible for a better user experience [34–36]. Currently, among the existing approaches, Riegler and Koltun stable view synthesis (SVS) approach [36] has shown excellent results and demonstrated photorealism in novel view synthesis, without using synthetic gaming engine 3D data, unlike [34]. SVS is indeed stable in rendering photorealistic images from novel viewpoints for large-scale scenes. Yet, it assumes MVS [39, 40] based dense 3D scene reconstruction and camera poses from COLMAP [39] are correct. The off-the-shelf algorithms used for 3D data acquisition and camera poses from images are, of course, popular, and to assume these algorithms could provide favorable 3D reconstruction and camera poses is not an outlandish assumption. Nonetheless, taking a step forward, in this paper, we argue that although choices made by SVS for obtaining geometric scaffold and camera poses in the pursuit of improving view synthesis is commendable, we can do better by making mindful use of fundamentals from multiple-view geometry [11, 12] and recent developments in deep-learning techniques for 3D computer vision problems.

To start with, we would like to emphasize that it is clearly unreasonable, especially in an outdoor setting, to assume that multi-view stereo (MVS) can provide accurate depth for all image pixels. It is natural that pixels with low disparity will not be reconstructed well using state-of-the-art MVS approaches [8, 39, 40, 46]. Even a precise selection of multiple view images with reasonable distance between them (assume good baseline for stereo) may not be helpful due to loss of common scene points visibility, foreshortening issue, etc. [43]. Such issues compel the practitioner to resort to post-processing steps for refining the MVS-based 3D geometry so that it can be helpful for rendering pipeline or neural-rendering network at train time.

Another critical component to view synthesis, which is often brushed aside in the literature is the accurate recovery of the camera poses. In neural view synthesis approaches such as [36], if the camera pose is wrong, the feature ag-

*Equal Contribution

[†]Corresponding Author (k.sur46@gmail.com)




Figure 1. **Qualitative comparison.** Our result compared to the popular SVS method [36] on the M60 scene of the tanks and temples dataset [18]. It is easy to observe that our approach can better render fine details in the scene. For this scene, the PSNR values for SVS [36] and our method are 19.1 and 20.8, respectively, demonstrating improved PSNR result.

gregation corresponding surface points could be misleading, providing inferior results. Therefore, we should have camera poses as accurate as possible. Unfortunately, despite the camera pose importance to this problem, discussion on improving camera pose is often ignored under the assumption that COLMAP [39] provide the best possible camera pose estimates for any possible scenarios. Practically speaking, this is generally not the case for outdoor scenes [5, 11, 15]. What is more surprising is that some recent benchmark datasets put COLMAP recovered poses as the ground-truth poses [33]. Hence, we want to get this out way upfront that a robust and better camera-poses estimates are vital for better modeling view synthesis problem.

From the above predication, it is apparent that a more mindful approach is required to make view synthesis approaches practically useful, automatic, and valuable for real-world application. To this end, we propose a principled and systematic approach that provides a better geometric scaffold and camera poses for reliable feature aggregation of the scene’s surface points, leading to improved novel-view synthesis results enabling superior photorealism.

In practice, we can have suitable initial camera poses from images using COLMAP. Yet, it must be refined further for improved image-feature aggregation corresponding to 3D surface points for neural rendering. It is well-studied in multiple-view geometry literature that we can improve and refine camera poses just from image key-point correspondences [10, 11]. Accordingly, we introduce a learning-based multiple motion averaging via graph neural network for camera pose recovery, where the pose graph is initialized using COLMAP poses for refinement.

Meanwhile, it is challenging to accurately recover the 3D geometry of scene points with low or nearly-zero disparity using MVS methods [12, 43]. Another bad news from the theoretical side is that a precise estimation of scene depth from a single image is unlikely¹, which is a correct statement and hard to argue. The good news is that advancements in deep-learning-based monocular depth prediction have led to some outstanding results in several practical

applications [24, 31]. Thus, at least practically, it seems possible to infer reliable monocular depth estimates up to scale. Using single image depth prediction, we can reason about the depth of scene points with low disparities. So, our proposed strategy is to use confidence based multiple-view stereo 3D that favours pixels with near-to-mid disparity and allows monocular depth estimates for the rest of the pixels. Overall depth is recovered after scaling all the scene depth appropriately using MVS reconstructed metric.

By encoding the image features via convolutional neural networks, we map the deep features to our estimated 3D geometric scaffold of the scene. Since we have better camera poses and scene reconstruction, we obtain and aggregate accurate feature vectors corresponding to each imaging view-rays—both from the camera to the surface point and from the surface point to viewing image pixels, giving us a feature tensor. We render the new image from the features tensor via a convolutional network and simultaneously refine the camera pose. In summary, our contributions are

- A systematic and principled approach for improved stable view synthesis enabling enhanced photorealism.
- The introduced approach exploits the complementary nature of MVS and monocular depth estimation to recover better 3D geometric scaffold of the scene. Meanwhile, the robust camera poses are recovered using graph neural network based multiple motion averaging.
- Our approach proposes an improved loss function to jointly optimize and refine for poses, neural image rendering, and scene representation showing superior results.

Our approach when tested on benchmark datasets such as Tank and Temples [18], FVS [35], Mip-NeRF 360 [1], and DTU [16] gives better image based rendering results with generally more than 1 dB PSNR gain (see Fig.1).

2. Background and Preliminaries

Our work integrates the best of novel view synthesis and multiple view geometry approaches in computer vision in a mindful way. Both novel view synthesis and MVS are classical problems in computer graphics and computer vision, with an exhaustive list of literature. Thus, we confine our re-

¹As several 3D scene points can have same image projection.

lated work discussion to the methods that is directly related to the proposed approach. The interested reader may refer to [4, 8, 12, 28, 38, 42, 44] for earlier and recent progress in these areas. Here, we briefly discuss relevant methods and the current state-of-the-art in neural image-based rendering.

(i) Uncalibrated Multi-View Stereo. Given the intrinsic camera calibration matrix, we can recover camera poses and the 3D structure of the scene using two or more images [8, 21, 22, 39, 40]. One popular and easy-to-use MVS framework is COLMAP [39], which includes several carefully crafted modules to estimate camera poses and sparse 3D structures from images. For camera pose estimation, it uses classical image key-point-based algorithms [13, 27, 45]. As is known that such methods can provide sub-optimal solutions and may not robustly handle outliers inherent to the unstructured images. Consequently, camera poses recovered using COLMAP can be unreliable, primarily for outdoor scenes. Moreover, since the 3D reconstruction via triangulation uses sparse key points, at best accurate semi-dense 3D reconstruction of the scene could be recovered. Still, many recent state-of-the-art methods in novel view synthesis heavily rely on it [25, 35, 36, 49].

(ii) Image-based rendering. Earlier image-based rendering approaches enabled novel view synthesis from images without any 3D scene data under some mild assumptions about the camera, and imaging [9, 23, 41]. Later with the development of multiple-view stereo approaches and RGBD sensing modalities, several works used some form of 3D data to improve image-based rendering. Popular works along this line includes [2, 6, 14, 20, 29]. In recent years, neural network-based methods have dominated this field and enabled data-driven approach to this problem with an outstanding level of photorealism [26, 34–36, 49].

Among all the current methods we tested, SVS [36] stands out in novel view synthesis from images taken from a freely moving camera. This paper proposes an approach closely related to the SVS pipeline. SVS involves learning a scene representation from images using its dense 3D structure and camera poses predicted via uncalibrated MVS approaches detailed previously. The recovered 3D structure of the scene is post-processed to have a refined geometric scaffold in mesh representation. Using refined scene 3D, it learns a feature representation for a set of images and then projects it along the direction of the target view to be rendered. Finally, SVS involves re-projection from feature to image space resulting in the synthesized image for the target view. Given \mathcal{S} , the 3D structure for a scene, feature representation network f_θ with θ representing the network parameters, and the rendering network \mathcal{G}_μ with parameters μ , the resulting image \mathcal{I}_r along the pose \mathbf{p} is rendered as

$$\mathcal{I}_r = \mathcal{G}_\mu(\phi_a(\mathcal{S}, \mathcal{I}_s, \mathbf{p}, f_\theta), \quad (1)$$

where, \mathcal{I}_s denotes the image set for a given seen. $\phi_a()$ ag-

gregates features for all the images along a given direction in \mathcal{I}_s predicted when passed through f_θ .

Contrary to the choices made by the SVS, in this work, we put forward an approach for better estimation of the overall scene 3D geometric scaffold—for both low and high-disparity pixels—and camera poses. Our approach enables a better aggregation of surface features allowing an improved level of realism in novel view synthesis.

3. Method

We begin with a discussion about formulating and estimating a better scene representation by carefully exploiting the complementary behavior of monocular depth and multi-view stereo approaches in 3D data acquisition. During our implementation, we use confidence measures for precise reasoning of estimated 3D coming from both the modalities and their accuracy. Next, we detail integrating a camera pose-refining optimization to estimate better camera poses for accurately projecting our 3D scene representation to the target view. Finally, the proposed joint optimization objective to refine structure representation and camera poses simultaneously is discussed.

3.1. Overview

Given a set of source images \mathcal{I}_s taken from a freely moving camera, the current state-of-the-art SVS [36] first estimates camera-poses \mathcal{P} and scene 3D structure \mathcal{S} . Next, it encodes the images via a CNN. Then, using the estimated camera, it maps the encoded features onto the scene 3D scaffold. For each point $\mathbf{x} \in \mathbb{R}^3$ on the scaffold, SVS query set of images in which \mathbf{x} is visible to obtain its corresponding feature vectors. This feature vector is then conditioned on the output view direction via a network to produce a new feature vector. Such a new feature vector is obtained for all the points on the scaffold to form a feature tensor, which is then decoded using CNN to synthesize the output image. Denoting the complete set of parameters for learning the feature representation and rendering as Θ , we can write SVS idea as

$$\Theta \sim \Psi(\Theta|\mathcal{S}, \mathcal{P}, \mathcal{I}_s) \circ \Psi(\mathcal{S}, \mathcal{P}|\mathcal{I}_s). \quad (2)$$

$\Psi(\cdot)$ symbolizes an abstract functional. Here, \mathcal{S} and \mathcal{P} are estimated using structure from motion and MVS [39, 40]. As mentioned before, we first aim to recover a much better 3D scene recovery by utilizing the complementary nature of the monocular depth and stereo depth. Furthermore, estimate improved camera poses using neural graph-based multiple motion averaging. We denote \mathcal{D}_s as monocular depth for each image in \mathcal{I}_s and $\mathcal{S}_R, \mathcal{P}_R$ as our scene 3D reconstruction and camera poses, respectively.

Once we estimate $\mathcal{S}_R, \mathcal{P}_R$, we jointly optimize for better neural scene representation and camera poses with the neural network parameters Θ using the proposed loss function

(Sec. 3.3). Our better scene representation and improved camera help in improved view-synthesis at test time. Our overall idea can be understood via the following equation:

$$\Theta, \mathcal{S}_R, \mathcal{P}_R \sim \Psi(\Theta, \mathcal{S}_R, \mathcal{P}_R | \mathcal{S}, \mathcal{P}, \mathcal{D}_s, \mathcal{I}_s) \quad (3)$$

3.1.1 3D Scene Features

Here, we describe feature aggregation for the output view obtained using the SVS pipeline [36] combined with the back-projected RGB-D features due to images and their monocular depth via a convolutional neural network (CNN). Given the output direction \mathbf{u} , for each $\mathbf{x} \in \mathbb{R}^3$ on the geometric scaffold along \mathbf{u} there is a subset of source images, in which \mathbf{x} is visible. Assuming the total number of images in this subset to be K , we denote the set of view directions corresponding to these images as $\{\mathbf{v}_k\}_{k=1}^K$.

(i) SVS Features. The source images in \mathcal{I}_s are first passed through a feature extraction CNN to obtain a feature tensor \mathcal{F}_k corresponding to k^{th} image. Denoting $f_k(\mathbf{x})$ as the feature corresponding to a point $\mathbf{x} \in \mathbb{R}^3$ in \mathcal{F}_k located at the projection (or bilinear interpolation) of \mathbf{x} on k^{th} image. To this end, SVS [36] proposed the following aggregation function (ϕ_α) to compute the feature for \mathbf{x} along \mathbf{u}

$$\phi_\alpha(\mathbf{u}, \{\{\mathbf{v}_k, f_k(\mathbf{x})\}\}) = \frac{1}{W} \sum_{k=1}^K \max(0, \mathbf{u}^T \mathbf{v}_k) f_k(\mathbf{x}), \quad (4)$$

where, $W = \sum_{k=1}^K \max(0, \mathbf{u}^T \mathbf{v}_k)$ is sum of all the weights.

(ii) Monocular Features. Given an image, we predict its depth with a per-pixel confidence score. For this, we use an existing monocular depth prediction network [32] pre-trained on Omnidata [7]. To predict a normalized depth prediction confidence score per pixel w_{ij} , we add another network on top of it (refer supplementary for details), which takes both images and depth as input.

$$\sum_i \sum_j w_{ij} \cdot \mathcal{L}(d_{ij}, \hat{d}_{ij}); \text{ where, } w_{ij} = \phi_w(\mathcal{I}_s, \mathcal{D}_s)_{ij}. \quad (5)$$

\mathcal{L} symbolizes the l_2 loss between the known pixel depth d_{ij} and the predicted depth \hat{d}_{ij} for (i, j) pixel in the 2D depth image. We call network ϕ_w as the ‘‘confidence-prediction head’’. We take the depths predicted by this network and fuse them with the source image using the Channel Exchange Network [47]. The fused information is then projected into N dimensional feature space using network \mathcal{F}_θ , which is a CNN-based architecture. To compute the feature corresponding to a pixel \mathbf{p} in the target view \mathbf{u} , we warp the features predicted for each source views \mathbf{v}_k to this target view but now using the monocular depth $d \in \mathcal{D}_s$ of the source view.

$$g_k(\mathbf{u}, \mathbf{p}) = f_k^m(\mathcal{W}(\mathbf{p}, \mathbf{u}, \mathbf{v}_k, d)) \quad (6)$$

where, g_k denotes the feature warping function, $f_k^m(p)$ corresponds to feature in the tensor $\mathcal{F}_\theta(\mathcal{I}_s^k, d)$ (corresponding to k^{th} source image and corresponding monocular depth) at pixel \mathbf{p} , and \mathcal{W} is the warping function. We now aggregate the warped features corresponding to each source image using a weighted sum based on confidence scores:

$$\phi_\alpha(\mathbf{u}, \{\{\mathbf{v}_k, f_k^m(\mathbf{p})\}\}) = \sum_{k=1}^K c_k(\mathbf{p}) f_k^m(\mathbf{p}) \quad (7)$$

here, c_k symbolizes the predicted depth confidence map in the k^{th} view.

(iii) Structure Feature Aggregation Per Scene Point.

Given we have features from two different modalities *i.e.*, monocular depth and MVS, we now propose an aggregated feature representation for the target view per scene point $h_f = h_\theta(h_m, h_s)$, where h_θ is a CNN-based neural network with parameters θ and h_m is the final representation for the monocular estimation and h_s for the stereo-estimation. We aim to attain maximal correlation from both of the input representations *i.e.*,

$$\lambda_1 \mathcal{C}(h_f, h_m) \cdot \mathcal{C}(h_f, h_s) + \lambda_2 \mathcal{C}(h_f, h_s) + \lambda_3 \mathcal{C}(h_f, h_m) \quad (8)$$

where, $\sum_{i=1}^3 \lambda_i = 1$; $\lambda_i > 0 \forall i \in \{1, 2, 3\}$ and $\mathcal{C}(\cdot)$ being the standard correlation function [3]. We set $\lambda_1 = 1$ for nearby points where confidence of monocular depth could be greater than a defined threshold (τ). $\lambda_2 = 1$ for nearby points where confidence of monocular depth is lower than a defined threshold and finally $\lambda_3 = 1$ for far points whose relative depth is greater than a pre-defined threshold (σ). Such a choice is made since stereo features might be less accurate in certain depth range due to low disparity or too close to the lens. The final aggregation h_f comprises of CNN networks ϕ_α and ϕ_β , parameterized by α and β , respectively. The transformed features from each of these modalities are then fed to a CNN network ϕ_ν with parameters ν . Overall aggregated features per scene point is represented as

$$h_f = h_\theta(h_m, h_s) = \phi_\nu(\phi_\alpha(h_m), \phi_\beta(h_s)). \quad (9)$$

3.1.2 Camera Pose Estimation

Given we have initially estimated pose-set \mathcal{P} , we use multiple motion averaging (MRA) to recover better camera poses. MRA is fast, efficient, and robust in camera pose estimation from noisy key-point correspondences and works well even for sparse image sets. MRA takes a global view of the camera recovery problem. A view-graph corresponding to the initial set of camera poses is constructed as a global representation, with each graph node as a pose. For pose optimization, MRA first performs rotation averaging to recover rotation and then solves for translation. Assume a

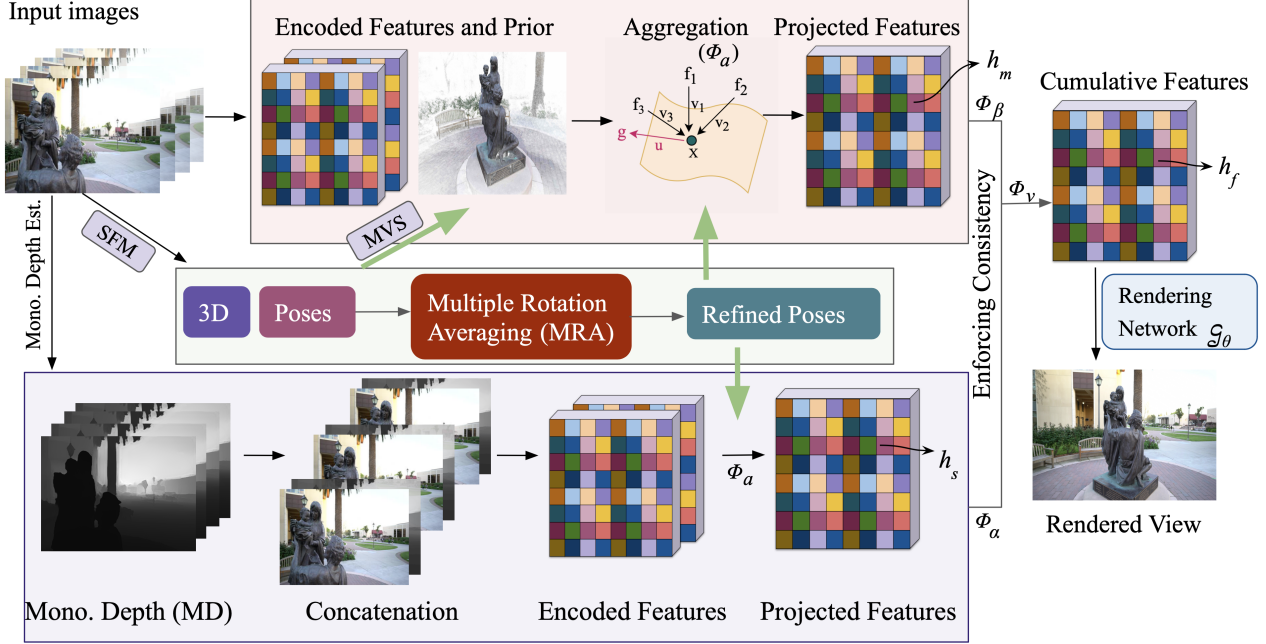


Figure 2. **Overview.** For a given set of input images, we first estimate the scene 3D structure and the initial value of camera poses using structure from motion (SFM) followed by MVS. We also estimate per-image depth using an existing Monocular Depth estimation model [7]. Next, we generate two sets of Projected Features for each target view: The first set (upper stream) by encoding the input images and then unprojecting the features onto the scaffold, followed by re-projection to the target view and aggregation. For the second set (lower stream), the images are first concatenated with their monocular depths (MD), and the encoded features are then projected using these depths, followed by aggregation. These two feature sets are then merged into Cumulative Features, and finally, the target view is rendered.

directed view-graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. A vertex $\mathcal{V}_j \in \mathcal{V}$ in this view graph corresponds to j^{th} camera absolute rotation R_j and $\mathcal{E}_{ij} \in \mathcal{E}$ corresponds to the relative orientation \tilde{R}_{ij} between view i and j . For our problem, relative orientations are given (noisy), and we aim to solve for absolute pose R_j robustly. Conventionally, in the presence of noise, the solution is obtained by solving the following optimization problem to satisfy compatibility criteria.

$$\operatorname{argmin}_{\{R_j\}} \sum_{\mathcal{E}_{ij} \in \mathcal{E}} \rho\left(\gamma(\tilde{R}_{ij}, R_j R_i^{-1})\right) \quad (10)$$

where, $\gamma(\cdot)$ denotes a metric on $SO(3)$ to measure distance between two rotation and $\rho(\cdot)$ is a robust l_1 norm. We use a graph-neural network to predict the robustified pose view-graph of the scene following the camera pose refining approach of NeuRoRA [30]. First, the camera poses are checked for outliers via a cyclic consistency in the pose-graph [11]. Then, a view graph is initialized based on these filtered poses. This view graph is then optimized using the camera pose-refining network of NeuRoRA [30].

3.2. Rendering novel views

The aggregated feature tensor \mathcal{F}^a along a direction \mathbf{u} comprises of the features h_f (Eq. 9) for each pixel along \mathbf{u} ,

obtained using 3D scene features (discussed in Sec. 3.1.1). This tensor is now projected to the image space function using \mathcal{G}_θ , a CNN-based network with parameters θ . It is quite possible that the regions in the target image may not be encountered in any of the source images. For those points, the values in feature tensor \mathcal{F} are set to $\mathbf{0}$, and thus, they require some inpainting. Denoting the set of test images by \mathcal{I}_t , the rendered image corresponding to the k^{th} image in the test set, $\hat{\mathcal{I}}_t^k$, is predicted based on the following equation:

$$\hat{\mathcal{I}}_t^k = \mathcal{G}_\theta(\mathcal{F}_k^a) \quad (11)$$

Similar to stable view synthesis, we parameterize the \mathcal{G}_θ function using a U-Net [37] style model which also deals with inpainting/hallucinating the newly discovered regions in the test image. Fig.(2) shows the complete overview of the proposed method.

3.3. Joint Optimization

Given \mathcal{I}_s , we train the model using our overall loss function \mathcal{L} comprising of two terms.

$$\mathcal{L} = \mathcal{L}_s + \mathcal{L}_p \quad (12)$$

(i) The **first loss term** \mathcal{L}_s encourages the network to learn better features corresponding to the scene point. The \mathcal{L}_s

	Truck			M60			Playground			Train		
	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow
NeRF++ [49]	21.8	0.31	0.81	17.6	0.43	0.73	21.9	0.39	0.79	17.6	0.48	0.68
FVS [35]	21.9	0.14	0.84	15.8	0.32	0.77	21.7	0.21	0.83	17.3	0.28	0.75
SC-NeRF [17]	22.3	0.29	0.82	18.4	0.40	0.76	22.4	0.35	0.83	18.2	0.42	0.73
Point-NeRF [48]	22.7	0.14	0.87	19.6	0.21	0.85	22.2	0.19	0.83	18.6	0.16	0.83
SVS [36]	22.9	0.12	0.88	19.1	0.22	0.83	22.9	0.17	0.86	17.9	0.19	0.81
Ours	24.1	0.12	0.90	20.8	0.20	0.89	23.9	0.14	0.90	20.1	0.13	0.88

Table 1. Rendered image quality comparison with current state-of-the-art methods in novel view synthesis on the popular Tanks and Temples dataset [18]. We use the popular metrics i.e., PSNR, LPIPS and SSIM for the comparison.

	Bike		Flowers		Pirate		Digger		Sandbox		Soccertable	
	SSIM \uparrow	LPIPS \downarrow	SSIM \uparrow	LPIPS \downarrow	SSIM \uparrow	LPIPS \downarrow	SSIM \uparrow	LPIPS \downarrow	SSIM \uparrow	LPIPS \downarrow	SSIM \uparrow	LPIPS \downarrow
NeRF++ [49]	0.71	0.27	0.80	0.31	0.71	0.43	0.65	0.35	0.84	0.24	0.87	0.21
FVS [35]	0.61	0.28	0.79	0.27	0.69	0.37	0.68	0.24	0.78	0.32	0.82	0.21
SVS [36]	0.74	0.22	0.84	0.21	0.75	0.32	0.77	0.18	0.85	0.20	0.91	0.15
Ours	0.79	0.19	0.88	0.18	0.80	0.29	0.81	0.16	0.91	0.17	0.94	0.13

Table 2. Quantitative comparison on FVS dataset [35]. We use the popular metrics, i.e., PSNR, LPIPS and SSIM for the comparison.

comprises to two objective function \mathcal{L}_{rgb} and \mathcal{L}_{corr} . The function \mathcal{L}_{rgb} measures the discrepancies between the rendering of an image in the source set and the actual image available in the set. This objective is used to train the structure estimation and rendering network parameters jointly. Whereas \mathcal{L}_{corr} is used for maximizing the correlation objective discussed in Eq.(8) to arrive at the optimal aggregated scene representation.

$$\mathcal{L}_s = \mathcal{L}_{rgb} + \mathcal{L}_{corr} \quad (13)$$

\mathcal{L}_{corr} takes the negative of the objective defined in Eq.(8). The structure network parameters corresponding to functions f_k and f_k^m (cf. Sec. 3.1.1) and is updated using \mathcal{L}_{corr} and \mathcal{L}_{rgb} . The rendering network parameters corresponding to \mathcal{G}_θ updated using only \mathcal{L}_{rgb} .

(ii) The **second loss term** \mathcal{L}_p corresponds to the camera pose refinement. We used the following loss \mathcal{L}_p to improve the camera pose estimation.

$$\mathcal{L}_p = \mathcal{L}_{mra} \quad (14)$$

As is known, the multiple motion averaging and the color rendering cost functions are directly impacted by the camera pose parameters. And therefore, the \mathcal{L}_{rgb} is inherently used as an intermediate objective in which the camera pose term is constant between the two different representations.

4. Experiments, Results and Ablations

Here, we discuss our experimental results and their comparison to relevant baselines. Later, we provide critical ablation analysis, a study of our camera pose estimation idea, and 3D structure recovery. Finally, the section concludes with a discussion of a few extensions of our approach².

²Please refer supplementary for the train-test setup, implementation details regarding hyperparameters, and architectures to reproduce our results.

Baselines. We compare our approach with the recently proposed novel view synthesis methods that work well in practice for real-world scenes. Our baseline list includes NeRF++ [49], SC-NeRF [17], PointNeRF [48], FVS [35] and SVS [36]. All baseline methods are trained on the source image set of the testing sequence. Our method’s train and test set are the same as SVS [36].

4.1. Dataset and Results

For evaluation, we used popular benchmark datasets comprising real-world and synthetic datasets. Namely, we used the Tanks and Temples [18], FVS [35], Mip-NeRF 360 [1], and the DTU [16] dataset. The first two datasets contain real-world scenes, whereas the last is an object-centric synthetic benchmark dataset. Although our approach mainly targets realistic scenes, for completeness, we performed and tabulated results on a synthetic object-based DTU dataset (see supplementary). Next, we discuss datasets and results.

(i) **Tanks and Temples Dataset.** It consists of images of large-scale real-world scenes taken from a freely moving camera, consisting of indoor and outdoor scenes. Unfortunately, we have no ground-truth poses; thus, estimated poses using COLMAP [39] are treated as pseudo-ground-truth poses. The dataset consists of a total of 21 scenes. Similar to the SVS setup, we use 15 of 21 scenes for training, 2 for validation, and the rest as test scenes. We also use a disjoint set of target views from the test scenes to evaluate all the methods. The source views in test scenes are used for scene-specific fine-tuning if required. Table 1 shows the statistical comparison of our approach with competing baselines. Our approach consistently outperforms the relevant baselines on this challenging dataset. Furthermore, if we zoom into the scene’s finer details, our approach clearly preserves details better than other approaches (see Fig. 3).

(ii) **FVS Dataset.** It comprises six real-world scenes. Each




Figure 3. Qualitative comparison on tanks and temples dataset. If we zoom into the scene details, our approach results show considerably less artifacts than the state-of-the-art methods enabling unparalleled level of realism in image-based rendering. Our PSNR values for the above two scenes are (20.1, 23.9). In comparison, SC-NeRF [17] and SVS [36] provide (17.3, 22.4) and (17.9, 22.9), respectively.

scene was recorded two or more times to completely differentiate the source image set to form the target image set. Table 2 summarises this dataset’s view synthesis quantitative results. We compare our method with the state-of-the-art NeRF++ [49], Free-View synthesis [35], and SVS [36] methods. We observe that FVS improves over NeRF++, whereas SVS results on this dataset do not consistently outperform other previous methods. On the contrary, our method consistently provides better results than prior art in all categories. The dataset’s sample images and its qualitative results are provided in the supplementary.

(iii) **Mip-NeRF 360 Dataset.** We further test our approach on this recently proposed dataset comprising unbounded scenes [1]. We use the train-test setup per scene proposed by the authors to fine-tune our network and compare it with Mip-NeRF 360 [1], NeRF++ [49], and SVS [36] baseline. Table 3 shows the results for this comparison. Clearly, our methodical take on this problem generalizes well, outperforming other baseline methods for most examples.

	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow
NeRF++ [49]	25.03	0.355	0.682
SVS [36]	25.28	0.218	0.783
PBNR [19]	23.55	0.262	0.722
Mip-NeRF 360 [1]	27.07	0.251	0.781
Ours	27.95	0.232	0.797

Table 3. Quantitative comparison of our approach with existing baselines on the Mip-NeRF 360 dataset.

4.2. Ablation Analysis

(i) **Effect of our camera-pose estimation approach.** We performed this experiment to show the benefit of using multiple motion averaging (MRA) instead of relying only on COLMAP poses and taking them for granted. For this experiment, we used the popular Lego dataset [25]. We first compute the COLMAP camera poses shown in Fig.4 (Left). Next, we use this initial camera pose set to perform MRA utilizing the view-graph optimization. Our recovered camera poses are shown in Fig.4 (Right). For clarity, we also show the ground-truth camera pose frustum and the error between the ground-truth pose and recovered pose for the




Figure 4. COLMAP [39] and our method’s camera poses on Lego dataset [25]. Here, we add some noise to the pairwise matched correspondences to simulate realistic scenarios. **Left:** Estimated camera pose from COLMAP and the ground truth pose. Here, black line shows the pose error. **Right:** Our recovered camera pose. It is easy to infer our approach robustness.

two respective cases. It is clear from Fig.(4) that a principled approach is required for better camera-pose recovery in view-synthesis problem. Furthermore, to show the benefit of our estimated camera pose in novel-view synthesis, we show our network training loss curve with and without our camera pose estimation approach in Fig.(5). For clarity, we also plot the SVS training curve. Firstly, our method shows better loss response due to the improved 3D geometric scaffold of the scene. Moreover, we improve on training loss further by introducing neural graph-based MRA, showing the benefit of our proposed loss function.




Figure 5. Training loss curve of our method with and without pose optimization compared to state-of-the-art SVS [36]

(ii) Effect of using MVS with monocular depth. We performed this ablation to demonstrate the effectiveness of our approach in recovering geometric reconstruction of the scene. Fig.(6) clearly shows the complementary nature of the monocular depth and stereo depth in 3D reconstruction from images. While MVS provides reliable depth results for near and mid-range points, monocular depth can deliver a good depth inference (up to scale) on the far points in the scene. By carefully integrating both modalities’ depth, we have better 3D information about the scene. Fig.(6) shows our results on a couple of scenes from Tanks and Temples dataset, demonstrating the suitability of our approach.




Figure 6. Comparison of depth predicted by monocular and stereo based method on the scenes from Tanks and Temples dataset.

5. Conclusion

The approach presented in this paper directly addresses the shortcomings of the currently popular methods in novel view synthesis. Our approach integrates concepts from the learning-based approaches and classical techniques in multiple-view geometry for solving novel-view synthesis. We demonstrate that by exploiting the (i) complementary nature of monocular depth estimation and multiple view stereo (MVS) in the 3D reconstruction of the scene, and (ii) usefulness of multiple rotation averaging in structure from motion, we can achieve better-rendering quality than other dominant approaches in novel view synthesis. We confirm the effectiveness of our approach via a comprehensive evaluation and quantitative comparison with baselines on several benchmark datasets. Although our work enabled improved photorealistic rendering, several exciting avenues exist for further improvement. One interesting future extension is fully automating our introduced view synthesis pipeline, *i.e.*, adding a learnable MVS framework to estimate intrinsic, extrinsic camera parameters with scene reconstruction for better novel view synthesis.

References

- [1] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022.
- [2] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstructured lumigraph rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 425–432, 2001.
- [3] Sarath Chandar, Mitesh M Khapra, Hugo Larochelle, and Balaraman Ravindran. Correlational neural networks. *Neural computation*, 28(2):257–285, 2016.
- [4] Yuan Chang and WANG Guo-Ping. A review on image-based rendering. *Virtual Reality & Intelligent Hardware*, 1(1):39–54, 2019.
- [5] Avishek Chatterjee and Venu Madhav Govindu. Robust relative rotation averaging. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):958–972, 2017.
- [6] Gaurav Chaurasia, Sylvain Duchene, Olga Sorkine-Hornung, and George Drettakis. Depth synthesis and local warps for plausible image-based navigation. *ACM Transactions on Graphics (TOG)*, 32(3):1–12, 2013.
- [7] Ainaz Eftekhari, Alexander Sax, Jitendra Malik, and Amir Zamir. Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10786–10796, 2021.
- [8] Yasutaka Furukawa, Carlos Hernández, et al. Multi-view stereo: A tutorial. *Foundations and Trends® in Computer Graphics and Vision*, 9(1-2):1–148, 2015.
- [9] Steven J Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F Cohen. The lumigraph. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 43–54, 1996.
- [10] Venu Madhav Govindu. Combining two-view constraints for motion estimation. In *CVPR*, volume 2. IEEE, 2001.
- [11] Richard Hartley, Jochen Trunpf, Yuchao Dai, and Hongdong Li. Rotation averaging. *International journal of computer vision*, 103(3):267–305, 2013.
- [12] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [13] Richard I Hartley. In defense of the eight-point algorithm. *IEEE Transactions on pattern analysis and machine intelligence*, 19(6):580–593, 1997.
- [14] Peter Hedman, Tobias Ritschel, George Drettakis, and Gabriel Brostow. Scalable inside-out image-based rendering. *ACM Transactions on Graphics (TOG)*, 35(6):1–11, 2016.
- [15] Nishant Jain, Suryansh Kumar, and Luc Van Gool. Robustifying the multi-scale representation of neural radiance fields. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*. BMVA Press, 2022.
- [16] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 406–413, 2014.
- [17] Yoonwoo Jeong, Seokjun Ahn, Christopher Choy, Anima Anandkumar, Minsu Cho, and Jaesik Park. Self-calibrating neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5846–5854, 2021.
- [18] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (TOG)*, 36(4):1–13, 2017.
- [19] Georgios Kopanas, Julien Philip, Thomas Leimkühler, and George Drettakis. Point-based neural rendering with per-view optimization. In *Computer Graphics Forum*, volume 40, pages 29–43. Wiley Online Library, 2021.
- [20] Johannes Kopf, Michael F Cohen, and Richard Szeliski. First-person hyper-lapse videos. *ACM Transactions on Graphics (TOG)*, 33(4):1–10, 2014.
- [21] Suryansh Kumar, Yuchao Dai, and Hongdong Li. Monocular dense 3d reconstruction of a complex dynamic scene from two perspective frames. In *Proceedings of the IEEE international conference on computer vision*, pages 4649–4657, 2017.
- [22] Suryansh Kumar, Yuchao Dai, and Hongdong Li. Superpixel soup: Monocular dense 3d reconstruction of a complex dynamic scene. *IEEE transactions on pattern analysis and machine intelligence*, 43(5):1705–1717, 2019.
- [23] Marc Levoy and Pat Hanrahan. Light field rendering. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 31–42, 1996.
- [24] Ce Liu, Suryansh Kumar, Shuhang Gu, Radu Timofte, and Luc Van Gool. Va-depthnet: A variational approach to single image depth prediction. In *The Eleventh International Conference on Learning Representations (ICLR)*, 2023.
- [25] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [26] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [27] David Nistér. An efficient solution to the five-point relative pose problem. *IEEE transactions on pattern analysis and machine intelligence*, 26(6):756–770, 2004.
- [28] Onur Özyeşil, Vladislav Voroninski, Ronen Basri, and Amit Singer. A survey of structure from motion*. *Acta Numerica*, 26:305–364, 2017.
- [29] Eric Penner and Li Zhang. Soft 3d reconstruction for view synthesis. *ACM Transactions on Graphics (TOG)*, 36(6):1–11, 2017.
- [30] Pulak Purkait, Tat-Jun Chin, and Ian Reid. Neurora: Neural robust rotation averaging. In *European Conference on Computer Vision*, pages 137–154. Springer, 2020.
- [31] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. *ICCV*, 2021.

- [32] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2020.
- [33] Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and David Novotny. Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10901–10911, 2021.
- [34] Stephan R Richter, Hassan Abu Al Haija, and Vladlen Koltun. Enhancing photorealism enhancement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [35] Gernot Riegler and Vladlen Koltun. Free view synthesis. In *European Conference on Computer Vision*, pages 623–640. Springer, 2020.
- [36] Gernot Riegler and Vladlen Koltun. Stable view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12216–12225, 2021.
- [37] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [38] Muhamad Risqi U Saputra, Andrew Markham, and Niki Trigoni. Visual slam and structure from motion in dynamic environments: A survey. *ACM Computing Surveys (CSUR)*, 51(2):1–36, 2018.
- [39] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [40] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *European conference on computer vision*, pages 501–518. Springer, 2016.
- [41] Steven M Seitz and Charles R Dyer. View morphing. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 21–30, 1996.
- [42] Harry Shum and Sing Bing Kang. Review of image-based rendering techniques. In *Visual Communications and Image Processing 2000*, volume 4067, pages 2–13. SPIE, 2000.
- [43] Richard Szeliski. *Computer vision: algorithms and applications*. Springer Nature, 2022.
- [44] Ayush Tewari et al. State of the art on neural rendering. In *Computer Graphics Forum*, volume 39, pages 701–727. Wiley Online Library, 2020.
- [45] Bill Triggs, Philip F. McLauchlan, Richard I. Hartley, and Andrew W. Fitzgibbon. Bundle adjustment - a modern synthesis. In *Proceedings of the International Workshop on Vision Algorithms: Theory and Practice, ICCV '99*, pages 298–372, London, UK, UK, 2000. Springer-Verlag.
- [46] Fangjinhua Wang, Silvano Galliani, Christoph Vogel, Pablo Speciale, and Marc Pollefeys. Patchmatchnet: Learned multi-view patchmatch stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14194–14203, 2021.
- [47] Yikai Wang, Wenbing Huang, Fuchun Sun, Tingyang Xu, Yu Rong, and Junzhou Huang. Deep multimodal fusion by channel exchanging. *Advances in Neural Information Processing Systems*, 33:4835–4845, 2020.
- [48] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5438–5448, 2022.
- [49] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020.