

Level-S²fM: Structure from Motion on Neural Level Set of Implicit Surfaces

Yuxi Xiao¹ Nan Xue^{1*} Tianfu Wu² Gui-Song Xia¹

¹ School of Computer Science, Wuhan University ² Department of ECE, NC State University

Project Page: <https://henry123-boy.github.io/level-s2fm/>

Abstract

This paper presents a neural incremental Structure-from-Motion (SfM) approach, Level-S²fM, which estimates the camera poses and scene geometry from a set of uncalibrated images by learning coordinate MLPs for the implicit surfaces and the radiance fields from the established keypoint correspondences. Our novel formulation poses some new challenges due to inevitable two-view and few-view configurations in the incremental SfM pipeline, which complicates the optimization of coordinate MLPs for volumetric neural rendering with unknown camera poses. Nevertheless, we demonstrate that the strong inductive basis conveying in the 2D correspondences is promising to tackle those challenges by exploiting the relationship between the ray sampling schemes. Based on this, we revisit the pipeline of incremental SfM and renew the key components, including two-view geometry initialization, the camera poses registration, the 3D points triangulation, and Bundle Adjustment, with a fresh perspective based on neural implicit surfaces. By unifying the scene geometry in small MLP networks through coordinate MLPs, our Level-S²fM treats the zero-level set of the implicit surface as an informative top-down regularization to manage the reconstructed 3D points, reject the outliers in correspondences via querying SDF, and refine the estimated geometries by NBA (Neural BA). Not only does our Level-S²fM lead to promising results on camera pose estimation and scene geometry reconstruction, but it also shows a promising way for neural implicit rendering without knowing camera extrinsic beforehand.

1. Introduction

Structure-from-Motion (SfM) is a fundamental 3D vision problem that aims at reconstructing 3D scenes and estimating the camera motions from a set of uncalibrated images. As a long-standing problem, there have been a tremendous of studies that are mostly established on the keypoint correspondences across viewpoints and the theoretical findings of Multi-View Geometry (MVG) [11], and

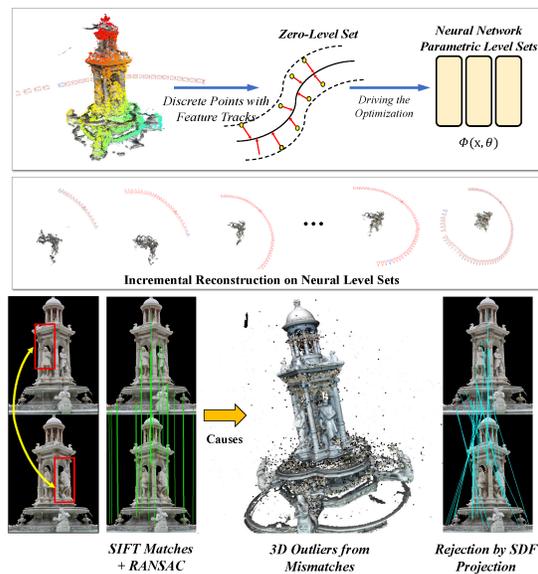


Figure 1. **SfM calculations on neural level sets.** We learn to do geometric calculations including Triangulation, PnP, and Bundle Adjustment above neural level sets, which easily help to reject the outliers in the matches especially in the texture repeated scenes. Also, due to the continuous surface priors of neural level sets, we achieve better pose estimation accuracy and our reconstructed points are sticking on the surface which are painted with color in the figure. While, there are a lot of outlier 3d points reconstructed by COLMAP [32] which are painted by black.

have formed three representative pipelines of Incremental SfM [32], Global SfM [4, 43], and Hybrid SfM [5].

In this paper, we focus on the incremental pipeline of SfM and we will use SfM to refer to the incremental SfM. Given an unordered image set, an SfM system initializes the computation by a pair of images that are with well-conditioned keypoint correspondences to yield an initial set of feature tracks, then incrementally adds new views one by one to estimate the camera pose from the 2D-3D point correspondences and update the feature tracks with new matches. Because the feature tracks are generated by grouping the putative 2D correspondences across viewpoints in bottom-up manners, they would be ineffective or inaccurate to represent holistic information of scenes. Accordingly,

*Corresponding author

Bundle Adjustment (BA) is necessary to jointly optimize the camera poses and 3D points in a top-down manner. The success of BA indicates that a global perspective is vital for accurate 3D reconstruction, however, their input feature tracks are the bottom-up cues without enough holistic constraints for the 3D scenes. To this end, we study to integrate the top-down information into the SfM system by proposing a novel Level-S²fM. Fig. 1 illustrates a representative case for the classic SfM systems that generate more flying 3D scene points, which can be addressed by our method.

Our Level-S²fM is inspired by the recently-emerged neural implicit surface that could manage all 3D scene points as the zero-level set of the signed distance function (SDF). Because the neural implicit surfaces can be parameterized by Multi-Layer Perceptrons (MLPs), it could be viewed as a kind of top-down information of 3D scenes and is of great potential for accurate 3D reconstruction. However, because both the 3D scene and camera poses are to be determined, it poses a challenging problem:

How can we optimize a neural SDF (or other neural fields such as NeRF) from only a set of uncalibrated images without any 3D information?

Most recently, the above problem was partially answered in BARF [18] and NeRF- - [42] that relaxed the requirement of optimizing Neural Radiance Fields [24] without knowing accurate camera poses, but they can only handle the unknown pose configurations in small-scale *face-forwarding* scenes. Moreover, when we confine the problem in the incremental SfM pipelines, it would be more challenging as we need to optimize the neural fields with only two overlapped images at the initialization stage. To this end, we found that the optimization of neural SDF can be accomplished by the 2D matches at the initialization stage, and facilitate the management of feature tracks by querying the 3D points and tracing the 2D keypoints in a holistic way.

As shown in Fig. 1, we define a neural network that parameterizes an SDF as the unified representation for the underdetermined 3D scene and accomplishes the computations of PnP for camera pose intersection, the 3D points triangulation as well as the geometry refinement on the parameterized SDF. In the initialization stage with a pair of overlapped images, Level-S²fM uses the differentiable sphere tracing algorithm [19] to attain the corresponding 3d points of the keypoints and calculate the reprojection error to drive the joint optimization. For the traced 3d points with small SDF values and 2D reprojection errors for its feature track, they are added into a dynamic point set and take the point set with feature tracks as the Lagrangian representation for the level sets. Because the pose estimation and the scene points reconstruction are sequentially estimated, the estimation errors will be accumulated. To this end, we present an NBA (*i.e.*, Neural Bundle Adjustment) that plays a similar

role as in Bundle Adjustment, but it optimizes the implicit surface and camera poses from the explicit flow of points by the energy function of the reprojection errors, which can be viewed as an evolutionary step between Lagrangian and Eulerian representations as discussed in [23].

In the experiments, we evaluate our Level-S²fM on a variety of scenes from the BlendedMVS [45], DTU [14], and ETH3D [34] datasets. On the BlendedMVS dataset, our proposed Level-S²fM clearly outperforms the state-of-the-art COLMAP [32] by significant margins. On the DTU and ETH3D datasets [14, 34], our method also obtains on-par performance with COLMAP for both camera pose estimation and dense surface reconstruction, which are all computed in one stage.

The contributions of this paper are in two folds:

- We present a novel neural SfM approach Level-S²fM, which formulates to optimize the coordinate MLP networks for implicit surface and radiance field and estimate the camera poses and scene geometry. To the best of our knowledge, our Level-S²fM is the first implicit neural SfM solution on the zero-level set of surfaces.
- From the perspective of neural implicit fields learning, we show that the challenging problems of two-view and few-view optimization of neural implicit fields can be addressed by exploiting the inductive biases conveyed in the 2D correspondences. Besides, our method presents a promising way for neural implicit rendering without knowing camera extrinsics beforehand.

2. Related Works

2.1. Structure from Motion

There has been a vast body of literature on Structure from Motion. Since an SfM system consists of many components, tremendous efforts have been devoted to improving the core components of SfM. In particular, the learning techniques were introduced in a variety of subproblems including image matching [7, 31], feature track mining and management [39], two-view 3D reconstruction [38, 44], relative and absolute camera pose estimation [13] and Bundle Adjustment [3, 37]. Those studies indicated that the learning paradigms are promising to improve the quality of 3D reconstruction. However, to the best of our knowledge, the learning paradigms are not fully equipped in SfM systems. One possible reason for such a fact is that the many learning approaches are designed in a supervised learning fashion, which remains some risks on the out-of-distribution samples. The self-supervised learning approaches [9, 22] in 3D vision alleviated the requirement of data annotations, however, they have not been fully exploited in the whole pipeline of SfM. In contrast to the aforementioned studies, in this paper, we are interested in integrating the learning

ability into the SfM system without incurring any external data annotations. From the perspective of system design in SfM, we verified that the strong inductive biases conveying in the 2D correspondences are promising and meaningful to drive the learning and optimization of SfM.

2.2. Neural Implicit Representation for 3D Scene

Recently, the advent of neural implicit fields [24, 28, 29, 40, 41, 46] have greatly advanced many 3D vision problems such as novel-view synthesis [1, 24] and surface reconstruction [28, 29, 40, 46] by learning to optimize the coordinate MLPs from a set of posed RGB images of which the key to success is that the inductive biases of 3D are exploited by the neural networks. However, when the camera poses are invalid, it is hard to optimize the coordinate MLPs for neural implicit fields. To remedy this, the state-of-the-art SfM system, COLMAP [32], is extensively used to compute the camera poses as a preprocessing step.

To train the neural field from unknown poses directly, recently, BARF [18] and NeRF [42] explored to jointly optimize the camera poses and neural fields by the volumetric rendering with promising results obtained in forward-facing scenes. BARF can also work in some scenes of highly overlapped and dense image collections with the initialized poses as inputs. This problem was also studied in the RGB-D SLAM systems [2, 36, 47], however, their works mainly rely on the known depth information and focus on the camera pose tracking by the neural implicit fields. Therefore, how to optimize implicit neural fields from only a set of uncalibrated images without any 3D information input is still a challenging and open problem.

In this paper, we study the unknown-pose neural fields optimization and SfM together and present a unified solution that simultaneously learns the implicit surfaces and radiance fields alongside the camera pose estimation and scene reconstruction from a set of images.

3. Preliminaries

In this section, we introduce the preliminaries on neural implicit surface rendering and the notations in SfM, which are all extensively used in our method.

3.1. Neural Implicit Surface Rendering

The volumetric rendering of neural implicit surface [46] aims at learning a signed distance function $d_\Omega : \mathbb{R}^3 \rightarrow \mathbb{R}$ by the volumetric rendering from a set of posed images and then extracting the zero-level set of ϕ as the reconstructed surface model of the image set. The state-of-the-art approach, VolSDF [46], integrates SDF representations with neural volume rendering via Laplacian distribution by

$$\sigma(\mathbf{x}) = \frac{1}{\beta} \Psi_\beta(-d_\Omega(\mathbf{x})), \quad (1)$$

where β is a learnable parameter in VolSDF [46]. Based on Eq. (1), the volume rendering equation renders a ray $\mathbf{x}(t)$ emanating from a camera position $\mathbf{o} \in \mathbb{R}^3$ in unit direction \mathbf{v} , defined by $\mathbf{x}(t) = \mathbf{o} + t\mathbf{v}$ by

$$I(\mathbf{o}, \mathbf{v}) = \int_0^\infty L(\mathbf{x}(t), \mathbf{n}(t), \mathbf{v}) \sigma(\mathbf{x}(t)) T(t) dt, \quad (2)$$

where $L(\mathbf{x}, \mathbf{n}, \mathbf{v})$ is the radiance field and $\mathbf{n}(t)$ is the normal direction of the point $\mathbf{x}(t)$ defined by $\mathbf{n}(t) = \nabla_{\mathbf{x}} d_\Omega(\mathbf{x}(t))$. In the learning of volume rendering, two coordinate MLP (Multi-Layer Perceptron) networks parameterize the SDF by $\phi(\mathbf{x}) = (d(\mathbf{x}), \mathbf{z}(\mathbf{x})) \in \mathbb{R}^{1+256}$ and the radiance field by $L_\psi(\mathbf{x}, \mathbf{n}, \mathbf{v}, \mathbf{z}) \in \mathbb{R}^3$, and train them by the color loss $\mathcal{L}_{\text{RGB}}(\phi, \psi, \beta)$ and the Eikonal loss $\mathcal{L}_{\text{Eik}}(\phi) = \mathbb{E}_{\mathbf{z}}(\|\nabla d(\mathbf{z})\| - 1)$.

In this paper, we use the equations (1) and (2) as the basic tools for Level-S²fM. To make the optimization of SDF and radiance networks easier, we set β as a small constant number and use the multi-resolution grid representations to avoid the potential of slow convergence and catastrophic forgetting since the scene scale is unknown and the original VolSDF [46] requires to normalize the known camera poses in a certain scale.

3.2. Ray Sampling and Sphere Tracing

Iterative Ray Sampling. In the implementation, the continuous form of Eq. (2) is approximated in

$$I(\mathbf{o}, \mathbf{v}) \approx \sum_{i=1}^{m-1} \hat{\tau}_i L(\mathbf{x}(t_i), \mathbf{n}(t_i), \mathbf{v}), \quad (3)$$

where $\{t_i\}_{i=1}^m$ is the discrete samples, $0 = t_1 < t_2 < \dots < t_m = M$, M is some large constant. $\hat{\tau}_i \approx \tau(s_i)\Delta s$ is the approximated PDF multiplied by the interval length. In VolSDF [46], $\{t_i\}_{i=1}^m$ is adaptively computed according to the opacity approximation error. Please move to [46] for its detail. In our method, we keep using this iterative sampling strategy when the rendering loss and the Eikonal loss is used. However, because the sampling set $\{t_i\}$ would be large, we do not use this strategy to compute the 3D points from 2D keypoints in our Level-S²fM and in turn to use the sphere tracing [10] as a faster way since our initial development of this work.

Sphere Tracing. Sphere tracing is a geometric method to render the depth from a signed distance function. Different from iterative ray sampling, sphere tracing is designed to hit the surface point along the ray $\mathbf{x}(t)$ with queries as few as possible. To make it clear, we use s_i to denote the ray stamp of the queried point $\mathbf{x}(s_i)$. With the queried point $\mathbf{x}(s_i)$, the next ray stamp s_{i+1} is computed by $s_{i+1} = \phi(\mathbf{x}(s_i))$. In our study, we sample at most $N_s = 20$ points with the stop criterion $|\phi(s_i)| < \varepsilon$, where ε is set to 0.002 in our experiment.

Remarks. Although both the iterative ray sampling [46] and sphere tracing [10] share the same target of computing the surface point along a ray, they have different behaviors in the neural implicit surface optimization. In detail, because VolSDF [46] aims at approximating the opacity by SDF, it updates the SDF network $\phi(\mathbf{x})$ by the rendering loss $\mathcal{L}_{\text{color}}$. As for sphere tracing, it is a geometric approach that only takes the SDF values into account for the computation. Such a difference is trivial to some extent, however, we found that their different focuses induce a loss function in our Level-S²fM to constraint the rendered depth values (or 3D points) for two-view initialization and 3D point triangulation.

3.3. Notations in SfM

Correspondence Search. Given the image set $\mathcal{I} = \{I_i | i = 1 \dots N_I\}$ for reconstruction, the keypoint features of the image I_i computed by SIFT [21] is denoted in $\mathcal{F}_i = \{(x_j, \mathbf{f}_j)\}$, where $x_j \in \mathbb{R}^2$ is the 2D coordinate and $\mathbf{f}_j \in \mathbb{R}^{128}$ is the feature descriptor of x_j . Based on the SIFT features, we follow the schema in COLMAP [32] to establish the feature correspondences across views, in which we first do the exhaustive matching for all possible image pairs and then use the geometric verification to filter out the non-overlapped image pairs. After this, the potentially overlapped image pairs are denoted in $\mathcal{C} = \{(I_a, I_b) | I_a, I_b \in \mathcal{I}\}$, and the keypoint correspondences in the pair (I_a, I_b) are denoted in the set $\mathcal{M}_{ab} = \{(x_k, \mathbf{f}_k), (x'_l, \mathbf{f}'_l) | (x_k, \mathbf{f}_k) \in \mathcal{F}_a, (x'_l, \mathbf{f}'_l) \in \mathcal{F}_b\}$. Finally, all the prepared correspondences are organized as the scene graph [32, 35], which stores images as the graph nodes and the overlapped image pairs as the graph edges. In our Level-S²fM, we use the established correspondences to drive the learning of MLPs, estimate the camera poses, and reconstruct a sparse point set of correspondences.

3D Scene Points and Feature Tracks. Because SfM is designated to simultaneously estimate the scene geometry from 2D correspondences, every successfully reconstructed 3D scene point is sourced from multiple 2D keypoint observations. To facilitate the representation, we denote the expected 3D point set in $\mathcal{X} = \{\mathbf{X}_k \in \mathbb{R}^3 | k = 1, \dots, N_{3d}\}$. For each point $\mathbf{X}_k \in \mathcal{X}$, if it is reconstructed from the 2D keypoint $x_j \in \mathcal{F}_i$, we denote such a relationship in a tuple (k, i, j) . $\mathcal{T} = \{(k, i, j)\}$ is the set of feature tracks.

4. The Proposed Level S²fM

In this section, we present the details of our Level-S²fM. As shown in Fig. 2, our method consists of three classical components including 1) the two-view geometry initialization, 2) the new frame registration, and 3) the new frame pose refinement, an implicit surface and a radiance field that are parameterized by neural networks. In what follows, we

will show how to solve the SfM problem by learning the implicit fields with 2D correspondences. We assume the intrinsic matrix K is known and fixed.

4.1. Two-view Initialization

We first select two good views $\{I_a, I_b\}$ for initialization from the scene graph and get their 2D matches $\mathcal{M}_{ab} = \{(x_k, \mathbf{f}_k), (x'_l, \mathbf{f}'_l) | (x_k, \mathbf{f}_k) \in \mathcal{F}_a, (x'_l, \mathbf{f}'_l) \in \mathcal{F}_b\}$. Based on the 2D matches \mathcal{M}_{ab} , we leverage the 5-point algorithm [27] and RANSAC to obtain the poses $P_a, P_b \in SE(3)$.

With the estimated camera poses P_a, P_b , it is straightforward to optimize the SDF network $\phi(\mathbf{x})$ and the radiance field network $L_\psi(\mathbf{x}, \mathbf{n}, \mathbf{v}, \mathbf{z})$ defined in Sec. 3 by minimizing the loss items \mathcal{L}_{RGB} and \mathcal{L}_{eik} as done in VolSDF [46]. However, it should be noted that the learning of volumetric surface rendering in such a way for the **two-view inputs** would trap into the local minimal by overfitting. To this end, we propose to use the differentiable sphere tracking [10, 19] for the corresponding rays in image I_a and I_b , which provides strong inductive biases for the optimization of networks.

Specifically, denoted by a pair of feature match (x_k, x'_l) in the image pair (I_a, I_b) , the sphere tracing obtains the surface point $\mathbf{X}_a^k = \mathbf{o}_a + \hat{t}_a \mathbf{d}_a$ from the SDF and $\mathbf{X}_b^l = \mathbf{o}_b + \hat{t}_b \mathbf{d}_b$, where $(\mathbf{o}_a, \mathbf{d}_a)$ is the ray of x_k , $(\mathbf{o}_b, \mathbf{d}_b)$ is the ray of x'_l . For the computation of \hat{t}_a and \hat{t}_b , please move to our supplementary materials. Ideally, the \mathbf{X}_a^k and \mathbf{X}_b^l should be as close as possible, therefore, we introduce a reprojection loss

$$\mathcal{L}_{\text{reproj}} = \frac{1}{2V} \sum (\|\hat{x}_k - x'_l\|_2 + \|\hat{x}'_l - x_k\|_2), \quad (4)$$

where V is the number of correspondences, $\hat{x}_k = \Pi(\mathbf{X}_a^k, K, P_b)$ and $\hat{x}'_l = \Pi(\mathbf{X}_b^l, K, P_a)$ are the projected 2D coordinates of the traced 3D points by the projection Π .

Considering the fact that the correspondences are sparse when the SDF network is not well optimized at some rays, the sparse sample points by sphere tracing on the SDF network may be either inaccurate or erroneous as shown in Fig. 3. Therefore, we use a depth consistency loss \mathcal{L}_{dc} to minimize the depth estimated by the sphere tracing and the volumetric rendering by

$$\mathcal{L}_{\text{dc}} = \frac{1}{B} \sum \|\hat{t}_i - \int_0^\infty T(t) \sigma(\mathbf{x}(t)) dt\|, \quad (5)$$

where the rays $\mathbf{x}(t)$ are randomly sampled from the images, and those rays are also used to compute the color loss \mathcal{L}_{RGB} . For the computation of Eikonal loss \mathcal{L}_{eik} , all the 3D points visited by sphere tracing and dense ray marching are used.

In summary, our two-view initialization of Level-S²fM computes the total loss $\mathcal{L}_{\text{total}}^{\text{init}}$ by

$$\mathcal{L}_{\text{total}}^{\text{init}} = \alpha_1 \mathcal{L}_{\text{reproj}} + \alpha_2 \mathcal{L}_{\text{eik}} + \alpha_3 \mathcal{L}_{\text{RGB}} + \alpha_4 \mathcal{L}_{\text{dc}}, \quad (6)$$

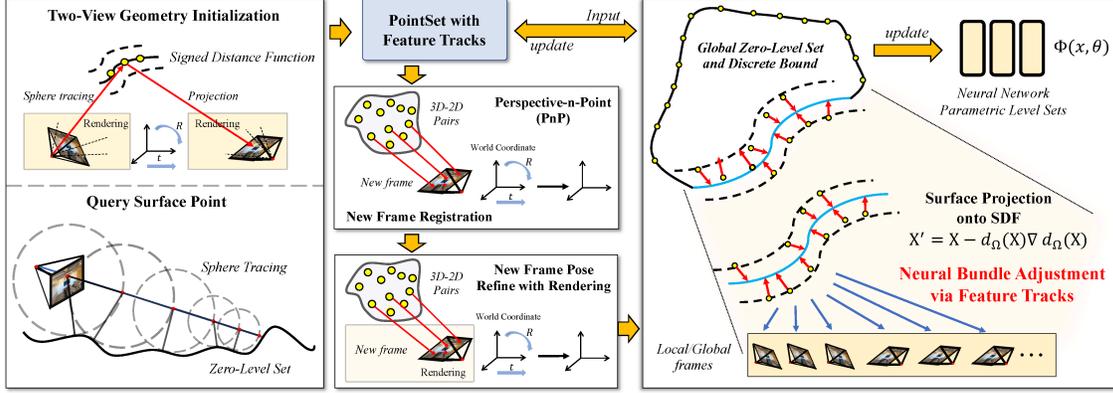


Figure 2. Overview of Level-S²fM.

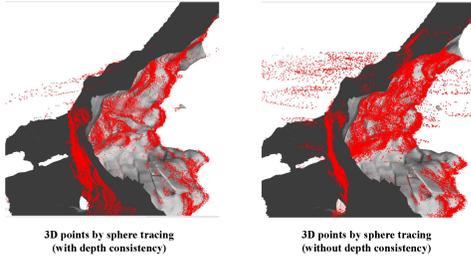


Figure 3. An illustrative comparison for the depth consistency loss \mathcal{L}_{dc} in two-view initialization. As it is shown, when we remove \mathcal{L}_{dc} , the 3D points traced from all putative 2D keypoints correspondences will contain more flying 3D points.

where $\alpha_1, \alpha_2, \alpha_3$, and α_4 are the hyperparameters and use ADAM optimizer to optimize the networks.

When the initialization is finished, we compute the two 3D points \mathbf{X} and \mathbf{X}' for each correspondence by sphere tracing for image I_a and I_b . For an accurate correspondence, $\|\mathbf{X} - \mathbf{X}'\|$ and their SDF values should be all small enough, thus providing a good criterion to check the putative matches to initialize the 3D point set \mathcal{X} and the feature track set \mathcal{T} for all the verified two-view correspondences.

4.2. New Frame Registration

For every newly added frame, we will first construct the 3D-2D correspondence from the existing pointset and its feature tracks. After that, we calculate a coarse pose of the new frame with the standard PnP algorithms [17], and then refine it with both the reprojection error and the rendering loss. The registration loss can be calculated as follow:

$$\mathcal{L}_{\text{regist}} = \beta_1 \mathcal{L}_{\text{reproj}} + \beta_2 \mathcal{L}_{\text{RGB}}, \quad (7)$$

where the β_1, β_2 are two hyper-parameters, and the $\mathcal{L}_{\text{reproj}}$ here are calculated by the 3D-2D correspondences.

In this optimization, the pose of the newly added frame, the SDF network, and the radiance field network are jointly

optimized. While during the changes in the pose and SDF, the original location in the pointset will maybe not be the right one on the surface. For this problem, we design a Neural Bundle Adjustment (NBA) strategy to dynamically update the pointset with respect to the SDF after the points triangulation and refinement in the next section. Therefore, we leave the details of NBA in Sec. 4.4.

4.3. Points Triangulation and Refinement

Once the pose of the newly added frame is obtained, we step into the next procedure of refining the retrieved 2D points from the point set \mathcal{X} and triangulating new 2D points into 3D space to extend \mathcal{X} . This problem was formulated in the classical SfM frameworks, however, they are suffering from the following issues:

- *The 2D Mismatches*: This issue could be alleviated by geometric verifications like RANSAC [8] or better 2D keypoint matching approaches, however, when encountering the symmetry structures or repeated texture regions, those efforts are hard to work efficiently.
- *Tiny Triangulation Angle*: This issue will lead to an ill-conditioned problem for points triangulation [12]. Therefore, the classical SfM approaches will directly discard those points to avoid the ill-conditioned problem configuration.

We address those issues by proposing an SDF-based triangulation. Similar to the two-view initialization in Sec. 4.1, we compute the 3D points for all the potential 2D keypoints in the first step. Then, for the 2D keypoints that have correspondences in the current feature track set \mathcal{T} , we use the tracing loss $\mathcal{L}_{\text{tracing}}$

$$\mathcal{L}_{\text{tracing}} = \frac{1}{V'} \sum_j \|\mathbf{X}_j^{st} - \mathbf{X}_j\|, \quad (8)$$

where $\mathbf{X}_j \in \mathcal{X}$ is the retrieved 3D point of the 2D keypoint in the current frame, V' is the number of retrieved 3D

points. This loss function acts as the similar role of \mathcal{L}_{dc} in two-view initialization. Without it, similar phenomenon like Fig. 3 will happen.

For the new 2D keypoints that are matched to the added images but without 3D information, both the reprojection loss similar to the two-view initialization and the tracing loss is used to yield the triangulation loss \mathcal{L}_{tri} by

$$\mathcal{L}_{tri} = \mathcal{L}_{reproj}^{\text{mask}} + \mathcal{L}_{tracing}, \quad (9)$$

where $\mathcal{L}_{reproj}^{\text{mask}}$ only considers the 2D correspondences of which their distance between the 2D projections of the traced 3D points in different views are smaller than a loose threshold (45 pixels in our implementation).

4.4. Neural Bundle Adjustment on Surfaces

Because the camera pose estimation and the points triangulation are separated, which will involve accumulative errors for pose estimation and triangulation, as well as the implicit networks. Motivated by the Bundle Adjustment that is extensively used in classical approaches, we present a Neural Bundle Adjustment (NBA) that jointly optimizes the estimated camera points, the 3D point set, and the implicit networks as a refinement step. To avoid costly computation, our NBA step finds the closest surface points to dynamically update those variables.

Denoted by the reconstructed 3D point set \mathcal{X} and the feature track \mathcal{T} , the camera poses $\mathcal{P} = \{P_1, \dots, P_K\}$ and the corresponding images $\{I_1, \dots, I_K\}$, as well as the networks $\phi(\mathbf{x})$ and L_ψ , in each step of NBA, we update the 3D point $\mathbf{X} \in \mathcal{X}$ by

$$\mathbf{X} \leftarrow \mathbf{X} - \phi(\mathbf{X})\nabla\phi(\mathbf{X}), \quad (10)$$

and then compute the reprojection loss according to the feature track \mathcal{T} to jointly optimize the ϕ the SDF network, \mathcal{P} the estimated camera poses, and \mathcal{X} the updated 3D point set. For the radiance network L_ψ , the rendering loss for randomly sampled rays is computed.

In our implementation, we leverage our NBA by three times, which we call the 1-frame NBA, local NBA, and global NBA. Because the rendering loss involves more rays, we only use it for the 1-frame NBA after the camera registration and point triangulation. In terms of local NBA, for the newly added view, only the related views with correspondences are considered. After running the 1-frame and local NBA schemes, we globally update all reconstructed views and the point set. By leveraging the backpropagation, all the mentioned variables are updated as the refinement.

5. Experiments

5.1. Implementation Details, Datasets, and Metrics

Implementation Details. In our implementation, we parameterize the SDF $\phi(\mathbf{x})$ by a multi-resolution features grid

and a two-layers MLP. To accelerate the computation, we follow InstantNGP [26] to use a hash table [25] for the feature grids. The radiance field L_ψ is also implemented in a multi-resolution feature grid and a three-layer MLP. Because our end task is the geometric 3D reconstruction, we use a high-resolution multi-scale feature grid for the SDF to ensure the accuracy of scene geometry but use a low-resolution feature grid to avoid the unnecessary computation cost for the radiance field. The specifications of the network architecture are given in supplementary material due to the limited space. All of these above are implemented in PyTorch [30], and we used the Adam [16] as the optimizer for the geometric calculations. For the 2D image matching and pose graph, we keep them the same with our baseline, COLMAP [32] for fair comparisons.

Datasets. Three datasets are used for our evaluation. Firstly, we use 5 representative scenes including the *LyingStatue*, *Stone*, *Fountain*, *Horse*, and *Statues* from the BlendedMVS dataset [45] in our evaluation because it provides accurate ground truth of camera poses and contains a number of challenging scenes for SfM. Secondly, the DTU dataset for the MVS task is also used. The five representative scenes (scans of 24, 37, 65, 110 and 114) are used in our experiments. Finally, we evaluate our proposed method on the five scenes from the challenging ETH3D [34] dataset.

Evaluation Metrics. In our evaluation, we use the Rotation error and ATE to quantitatively benchmark the pose accuracy, which simply depicts the difference between the ground truth and the aligned pose. During our evaluation, we used the provided API of Reconstruction Align in COLMAP [32] to do that. In terms of the reconstructed scene geometry, we use accuracy (Acc) and the precision (Prec) rate to evaluate the accuracy of our recovered 3D points and Chamfer- $l1$ distance to depict the accuracy of the reconstructed surface. Detailed definitions of these evaluation metrics are given in the supplementary material.

5.2. Results on the BlendedMVS Dataset

Tab. 1 reports the quantitative evaluation results for the two versions of Level-S²fM and COLMAP [32]. The full version of Level-S²fM used all the mentioned components while the wo/render version removed the rendering loss for optimization. As it is reported, our Level-S²fM (full) consistently outperforms COLMAP [32] for camera pose estimation and sparse 3D point cloud reconstruction. It also reveals that rendering losses are required.

In detail, our Level-S²fM (full) averagely reduced the estimation error from 1.54° by COLMAP [32] to 0.86°, **obtaining a relative improvement of 55.84%**. For the translation error, our Level-S²fM (full) decreases the error from 3.54 cm to 3.36 cm. For the sparse 3D point cloud reconstruction, the ACC metric is reduced from 3.16 to 2.25 for the full model and 2.63 for the wo/render version.

Scenes	Camera Pose Evaluation						Points Cloud Results Evaluation		
	Rotation ($^{\circ}$) \downarrow			Translation (cm) \downarrow			Acc(cm) \downarrow , Prec(≤ 3.5 cm) \uparrow		
	COLMAP [32]	Level-S ² fM (full)	Level-S ² fM (wo/render)	COLMAP [32]	Level-S ² fM (full)	Level-S ² fM (wo/render)	COLMAP [32]	Level-S ² fM (full)	Level-S ² fM (wo/render)
LyingStatue	1.20	1.12	1.31	0.89	2.18	2.67	Acc:0.83 Prec:0.99	Acc:1.35 Prec:0.99	Acc:1.76 Prec:0.98
Stone	0.63	0.31	0.54	6.28	5.51	9.17	Acc:5.44 Prec:0.66	Acc:3.61 Prec:0.68	Acc:5.04 Prec:0.64
Fountain	4.34	1.65	2.32	7.41	2.87	4.11	Acc:1.91 Prec:0.91	Acc:1.13 Prec:0.98	Acc:1.14 Prec:0.96
Horse	0.33	0.94	0.92	1.18	5.71	7.58	Acc:3.99 Prec:0.86	Acc:4.18 Prec:0.71	Acc:4.18 Prec:0.70
Statues	1.21	0.36	0.44	1.98	0.56	0.62	Acc:1.31 Prec:0.98	Acc:0.95 Prec:0.99	Acc:1.02 Prec:0.99
Mean	1.54	0.86	1.11	3.54	3.36	4.83	Acc:3.16 Prec:0.85	Acc: 2.25 Prec: 0.87	Acc:2.63 Prec:0.85

Table 1. **Quantitative results on the BlendedMVS dataset.** For our Level-S²fM, we report the results by full version and an wo/render version that removes the rendering loss during optimization.

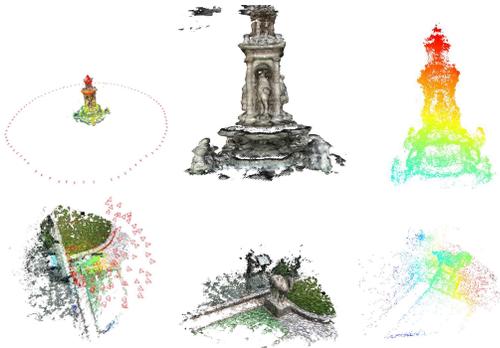


Figure 4. The reconstructed meshes, point clouds and camera poses for the *Fountain* and *Horse* scenes on the BlendedMVS dataset by our Level-S²fM (full). In the first column, the reconstructed scene geometry and the camera poses are shown together. For the 3D models, we show the different views of the sparse 3D points during the training and the textured meshes refused from the zero-level set surface.

Pose Source	LyingStatue	Stone	Fountain	Horse	Statues	Mean
COLMAP	29.1	28.4	25.4	23.5	29.2	27.12
Level-S ² fM (Ours)	29.5	28.9	27.1	23.6	30.2	27.86
GT	29.5	29.4	27.4	24.1	31.1	28.3

Table 2. **Novel View Synthesis Comparison.** The PSNR is used to compare the camera poses computed by COLMAP, Level-S²fM and the GT poses on the BlendedMVS dataset.

Fig. 4 shows the reconstruction results by our method.

Apart from the direct evaluation of the SfM results on the BlendedMVS dataset, we further compare the camera poses estimation results for different methods by training the NGP [26] (a fast version of NeRF [24]) to compare the performance of novel view synthesis in Tab. 2. As it is reported, the rendered images by our camera poses are consistently better than the ones by COLMAP poses.

5.3. Results on the DTU Dataset

We conducted the evaluation on the DTU to illustrate the promising future of our Level-S²fM to unify the pose estimation, dense reconstruction, and novel view synthesis problems in one stage. For the comparison to COLMAP, we use their built-in PatchMatch MVS [33] functionality to obtain the dense surface points and then leverage its default surface reconstruction method (*i.e.*, Poisson surface [15])

Scan	COLMAP [32]			Level-S ² fM (Ours)		
	Chamfer- ℓ_1	Rot. Err.	Trans. Err.	Chamfer- ℓ_1	Rot. Err.	Trans. Err.
24	2.176	0.38	2.87	2.442	0.81	4.60
37	3.837	0.41	4.86	3.023	0.31	4.29
65	4.394	0.45	4.23	3.190	0.74	5.81
110	3.389	0.65	6.36	5.902	0.82	6.82
114	3.577	0.35	3.58	2.092	0.14	1.85
Mean	3.330	0.448	4.38	3.474	0.564	4.67

Table 3. **Quantitative results on DTU dataset.** The Chamfer- ℓ_1 distance of the dense reconstruction results and as the errors of rotation and translation for camera pose estimation, are compared for COLMAP and our Level-S²fM. The unit of Chamfer- ℓ_1 and Translation errors are in millimetres.

Scene	Detector & Matcher	courtyard	relief	door	terrace2	facade
COLMAP [32]	SIFT [21]	0.10*/0.016m	0.10*/0.003m	0.16*/0.002m	0.14*/0.002m	0.06*/0.016m
Level-S ² fM (Ours)	SIFT [21]	0.21*/0.047m	0.09*/0.003m	0.19*/0.006m	0.13*/0.003m	0.12*/0.039m
COLMAP [32]	SP [7]+SG [31]	-	0.36*/0.007m	0.37*/0.003m	0.16*/0.003m	0.04*/0.014m
Level-S ² fM	SP [7]+SG [31]	-	0.42*/0.003m	0.34*/0.007m	0.12*/0.002m	0.10*/0.051m

Table 4. **Quantitative results** of pose estimation on the five scenes in ETH3D dataset for COLMAP [32] and our proposed Level-S²fM by using different keypoint detector and matcher.

to obtain the mesh model. For our Level-S²fM, we use the MarchingCubes [20] to extract the mesh models from the zero-level set of the implicit surface. The quantitative evaluation results are shown in Tab. 3. In this dataset, our Level-S²fM obtains on-par performance with COLMAP.

5.4. Results on the ETH3D Dataset

We test our method on a more challenging dataset, ETH3D [34], which includes both sparse view collections for multi-scale outdoor and indoor scenes. To show the influence of different keypoint detection and matching algorithms for our method, we additionally make a comparison with SuperPoint (SP) [7] for detection and SuperGlue (SG) [31] for keypoint matching. As reported in Tab. 4, our method achieves comparable results with COLMAP [32]. However, we observe that our method gets slightly inferior results in some large-scale outdoor scenes, because of the limited representative capability of a single network for a large-scale scene.

5.5. Ablation Study

In this section, we elaborate on why and how the SDF-based Triangulation (short in *SDF-Tri*) and NBA work in

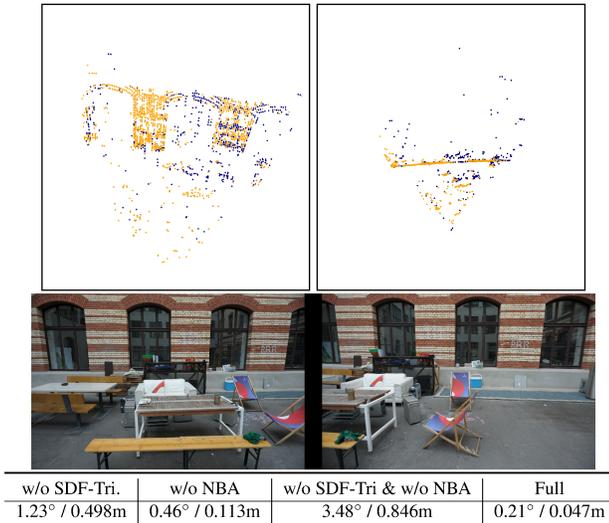


Figure 5. Two-views Triangulated Point Clouds by [Traditional triangulation](#) and [SDF-Based Triangulation](#), courtyard@ETH3D.

our system. We first conduct the two-view triangulation with SDF-Tri and the traditional method respectively. Fig. 5 shows SDF-Tri can easily filter out the incorrect triangulation from wrong matches. One explanation is that the neural network distills inliers in a continuous zero-level set of surfaces (e.g., the planes in Fig. 5). Therefore, the triangulated outliers (the blue points flying off the planes) can be easily detected and filtered by their large SDF values. Similarly, NBA also benefits from the global level sets of surfaces, which average the errors among the inliers triangulated points and differs the outliers by their large SDF values. The quantitative results for SDF-Tri and NBA are shown at the bottom of Fig. 5.

5.6. Limitations of Level-S²fM

In order to explore the clear boundary of Level-S²fM and point out the potential future development, we discuss the limitation of Level-S²fM on the most typical indoor dataset, scannet [6]. In the Scannet [6], there are a lot of challenges including blurry images, and textureless areas. Because of the less texture, the SIFT-based keypoint correspondences may contain a large portion of outliers or insufficient matches. Meanwhile, the blur in images will also influence the accuracy of the 2d matches. Therefore, most SfM easily fails on this dataset. Our method is also limited by this because of SIFT matches.

To make the discussion clear, we run four scenes of Scannet [6] that were used for NICESLAM [47]. For the image sequence of each scene, the input image set for SfM is constructed by sampling for every 10 frames. Tab. 5 reports the pose accuracy of COLMAP and our Level-S²fM for results by adding 60 frames, 120 frames, and all the frames, which are concatenated by the “slash”. As it is reported, our method usually performs well in the first 60 frames but its pose estimation accuracy suddenly decreases

Scene ID		0000	0059	0169	0207
IMAP [36]	rot ² ↓	–	–	–	–
	trans. (cm) ↓	197.1	18.9	96.4	28.7
NICESLAM [47]	rot ² ↓	–	–	–	–
	trans. (cm) ↓	11.3	12.0	12.0	12.8
COLMAP [32]	rot ² ↓	1.47/2.71/2.71	1.40/2.52/2.66	2.64/1.84/2.52	failed
	trans. (cm) ↓	8.20/16.8/14.1	7.20/10.8/11.5	10.6/9.22/10.6	–
Level-S²fM	rot ² ↓	1.44/2.23/1.98	1.50/2.59/3.72	2.94/9.09/14.3	failed
	full trans. (cm) ↓	6.50/11.8/26.7	8.20/32.9/35.7	6.35/14.3/44.8	–

Table 5. **Quantitative results** of pose estimation on ScanNet [6]. For the COLMAP and Level-S²fM, we report their pose accuracy metrics when the 60/120/all frames are registered. For the last scene (*i.e.*, 0207), both COLMAP and our method failed.

when adding some new frames. The textureless matches cause the matches very sparse and therefore hard to give a good registration of images. Meanwhile, since the radiance field learning is also challenging in the Scannet dataset, the bad initialized pose can not be refined well by the rendering loss. All of these limitations are basically from sparse 2D image matches. Besides, we observed that the ADAM optimizer will make the optimization of camera poses and scene points unstable, which would also affect the final results.

6. Conclusion

This paper studies the longstanding problem of Structure-from-Motion by exploring and exploiting several important yet challenging issues including the two-view neural rendering in the initialization stage and few-view neural rendering in the early camera registration stage of incremental SfM for integrating the recent advances of neural implicit field learning into an SfM pipeline. We show that although the few-view neural rendering problem is challenging enough, it can be tackled by the 2D correspondences as they convey strong inductive biases for 3D scenes. Based on this, we present the first neural SfM solution that renews several key components of two-view geometry initialization, camera pose registration, and triangulation, as well as the Bundle Adjustment problem with neural implicit fields. In the experiments, we show that Level-S²fM outperforms the traditional SfM pipeline and set a new state-of-the-art for 3D reconstruction on the BlendedMVS dataset. We believe that our study will encourage the 3D vision community to rethink and reformulate Structure-from-Motion with learning-based new findings.

Acknowledgments. This work was supported by the National Nature Science Foundation of China under grants 62101390 and U22B2011. This work was also supported by Wuhan University-Huawei Geoinformatics Innovation Laboratory. Y. Xiao would love to thank Sida Peng for the helpful discussions. T. Wu was supported by NSF IIS-1909644, IIS-1822477, CMMI-2024688, and IUSE-2013451. We are also grateful for the constructive comments by anonymous reviewers and area chairs. The views presented in this paper are those of the authors and should not be interpreted as representing any funding agencies.

References

- [1] Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *IEEE International Conference on Computer Vision (ICCV)*, pages 5835–5844, 2021. [3](#)
- [2] Ronald Clark. Volumetric bundle adjustment for online photorealistic scene capture. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6114–6122, 2022. [3](#)
- [3] Ronald Clark, Michael Bloesch, Jan Czarnowski, Stefan Leutenegger, and Andrew J. Davison. Learning to solve nonlinear least squares for monocular stereo. In *European Conference on Computer Vision (ECCV)*, volume 11212, pages 291–306, 2018. [2](#)
- [4] David J. Crandall, Andrew Owens, Noah Snavely, and Daniel P. Huttenlocher. Sfm with mrfs: Discrete-continuous optimization for large-scale structure from motion. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 35(12):2841–2853, 2013. [1](#)
- [5] Hainan Cui, Xiang Gao, Shuhan Shen, and Zhanyi Hu. Hsfm: Hybrid structure-from-motion. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2393–2402, 2017. [1](#)
- [6] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas A. Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3d reconstructions of indoor scenes. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2432–2443, 2017. [8](#)
- [7] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 224–236, 2018. [2, 7](#)
- [8] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. [5](#)
- [9] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J. Brostow. Digging into self-supervised monocular depth estimation. In *IEEE International Conference on Computer Vision (ICCV)*, pages 3827–3837, 2019. [2](#)
- [10] John C. Hart. Sphere tracing: a geometric method for the antialiased ray tracing of implicit surfaces. *The Visual Computer*, 12(10):527–545, 1996. [3, 4](#)
- [11] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge University Press, 2003. [1](#)
- [12] Richard I Hartley and Peter Sturm. Triangulation. *Computer Vision and Image Understanding (CVIU)*, 68(2):146–157, 1997. [5](#)
- [13] Petr Hruby, Timothy Duff, Anton Leykin, and Tomáš Pajdla. Learning to solve hard minimal problems. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5522–5532, 2022. [2](#)
- [14] Rasmus Ramsbøl Jensen, Anders Lindbjerg Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 406–413, 2014. [2](#)
- [15] Michael M. Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Symposium on Geometry Processing*, volume 256, pages 61–70, 2006. [7](#)
- [16] Diederik P Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*, 2015. [6](#)
- [17] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. EpnP: An accurate $O(n)$ solution to the pnp problem. *International Journal of Computer Vision (IJCV)*, 81(2):155–166, 2009. [5](#)
- [18] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. BARF: bundle-adjusting neural radiance fields. In *IEEE International Conference on Computer Vision (ICCV)*, pages 5721–5731, 2021. [2, 3](#)
- [19] Shaohui Liu, Yinda Zhang, Songyou Peng, Boxin Shi, Marc Pollefeys, and Zhaopeng Cui. DIST: rendering deep implicit signed distance function with differentiable sphere tracing. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2016–2025, 2020. [2, 4](#)
- [20] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Proceedings of SIGGRAPH*, pages 163–169, 1987. [7](#)
- [21] David G Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision (IJCV)*, 60(2):91–110, 2004. [4, 7](#)
- [22] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised Learning of Depth and Ego-Motion From Monocular Video Using 3D Geometric Constraints. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5667–5675, 2018. [2](#)
- [23] Ishit Mehta, Manmohan Chandraker, and Ravi Ramamoorthi. A level set theory for neural implicit evolution under explicit flows. In *European Conference on Computer Vision (ECCV)*, 2022. [2](#)
- [24] 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. In *European Conference on Computer Vision (ECCV)*, volume 12346, pages 405–421, 2020. [2, 3, 7](#)
- [25] Thomas Müller. tiny-cuda-nn, 4 2021. [6](#)
- [26] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *ACM Trans. on Graphics (TOG)*, 41(4):102:1–102:15, 2022. [6, 7](#)
- [27] David Nistér. An Efficient Solution to the Five-Point Relative Pose Problem. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 26(6):756–777, 2004. [4](#)
- [28] Michael Oechsle, Songyou Peng, and Andreas Geiger. UNISURF: unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *IEEE International Conference on Computer Vision (ICCV)*, pages 5569–5579, 2021. [3](#)
- [29] Jeong Joon Park, Peter Florence, Julian Straub, Richard A. Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation.

- In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 165–174, 2019. [3](#)
- [30] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 8024–8035, 2019. [6](#)
- [31] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4937–4946, 2020. [2, 7](#)
- [32] Johannes L Schönberger and Jan-Michael Frahm. Structure-from-motion Revisited. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4104–4113, 2016. [1, 2, 3, 4, 6, 7, 8](#)
- [33] 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 (ECCV)*, volume 9907, pages 501–518, 2016. [7](#)
- [34] Thomas Schöps, Torsten Sattler, and Marc Pollefeys. BAD SLAM: bundle adjusted direct RGB-D SLAM. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 134–144, 2019. [2, 6, 7](#)
- [35] Noah Snavely, Steven M. Seitz, and Richard Szeliski. Skeletal graphs for efficient structure from motion. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008. [4](#)
- [36] Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew J. Davison. imap: Implicit mapping and positioning in real-time. In *IEEE International Conference on Computer Vision (ICCV)*, pages 6209–6218, 2021. [3, 8](#)
- [37] Chengzhou Tang and Ping Tan. BA-Net: Dense Bundle Adjustment Networks. In *International Conference on Learning Representations (ICLR)*, 2019. [2](#)
- [38] Jianyuan Wang, Yiran Zhong, Yuchao Dai, Stan Birchfield, Kaihao Zhang, Nikolai Smolyanskiy, and Hongdong Li. Deep two-view structure-from-motion revisited. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8953–8962, 2021. [2](#)
- [39] Lei Wang, Linlin Ge, Shan Luo, Zihan Yan, Zhaopeng Cui, and Jieqing Feng. Tc-sfm: Robust track-community-based structure-from-motion. *CoRR*, abs/2206.05866, 2022. [2](#)
- [40] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 27171–27183, 2021. [3](#)
- [41] Yiqun Wang, Ivan Skorokhodov, and Peter Wonka. Improved surface reconstruction using high-frequency details. *CoRR*, abs/2206.07850, 2022. [3](#)
- [42] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf-: Neural radiance fields without known camera parameters. *CoRR*, abs/2102.07064, 2021. [2, 3](#)
- [43] Kyle Wilson and Noah Snavely. Robust global translations with Idsfm. In *European Conference on Computer Vision (ECCV)*, volume 8691, pages 61–75, 2014. [1](#)
- [44] Yuxi Xiao, Li Li, Xiaodi Li, and Jian Yao. Deepmle: A robust deep maximum likelihood estimator for two-view structure from motion. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10643–10650, 2022. [2](#)
- [45] Yao Yao, Zixin Luo, Shiwei Li, Jingyang Zhang, Yufan Ren, Lei Zhou, Tian Fang, and Long Quan. Blendedmvs: A large-scale dataset for generalized multi-view stereo networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1787–1796, 2020. [2, 6](#)
- [46] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 4805–4815, 2021. [3, 4](#)
- [47] Zihan Zhu, Songyou Peng, Viktor Larsson, Weiwei Xu, Hujun Bao, Zhaopeng Cui, Martin R. Oswald, and Marc Pollefeys. NICE-SLAM: neural implicit scalable encoding for SLAM. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12776–12786, 2022. [3, 8](#)