

PNeRFLoc: Visual Localization with Point-Based Neural Radiance Fields

Boming Zhao¹, Luwei Yang², Mao Mao¹, Hujun Bao¹, Zhaopeng Cui^{1*}

¹State Key Lab of CAD & CG, Zhejiang University

²Simon Fraser University

bmzhao@zju.edu.cn, mluweiyang@outlook.com, maomao6006@gmail.com, zhpcui@gmail.com, bao@cad.zju.edu.cn

Abstract

Due to the ability to synthesize high-quality novel views, Neural Radiance Fields (NeRF) has been recently exploited to improve visual localization in a known environment. However, the existing methods mostly utilize NeRF for data augmentation to improve the regression model training, and their performances on novel viewpoints and appearances are still limited due to the lack of geometric constraints. In this paper, we propose a novel visual localization framework, *i.e.*, PNeRFLoc, based on a unified point-based representation. On one hand, PNeRFLoc supports the initial pose estimation by matching 2D and 3D feature points as traditional structure-based methods; on the other hand, it also enables pose refinement with novel view synthesis using rendering-based optimization. Specifically, we propose a novel feature adaption module to close the gaps between the features for visual localization and neural rendering. To improve the efficacy and efficiency of neural rendering-based optimization, we also developed an efficient rendering-based framework with a warping loss function. Extensive experiments demonstrate that PNeRFLoc performs the best on the synthetic dataset when the 3D NeRF model can be well learned, and significantly outperforms all the NeRF-boosted localization methods with on-par SOTA performance on the real-world benchmark localization datasets. The code and supplementary material are available on the project webpage: <https://zju3dv.github.io/PNeRFLoc/>.

Introduction

Visual localization is a fundamental task in computer vision that aims to determine the precise position and orientation of a camera in a known scene based on the visual input, and it has widespread applications in areas such as robot navigation, augmented reality, virtual reality, *etc.* Traditional structure-based localization, as the mainstream solution for visual localization, has advantages such as scene agnosticism, robustness, and high precision. These methods (Brachmann and Rother 2021; Sarlin et al. 2019) require computing and storing a global map consisting of 3D point locations and try to find the correspondences between 2D feature points extracted in the query image and 3D points in the reconstructed scene and use a Perspective-n-Point (PnP) solver (Haralick et al. 1994; Bujnak, Kukelova, and Pajdla 2008)

in a RANSAC loop (Fischler and Bolles 1981; Chum and Matas 2008) to compute the camera poses. Besides hand-crafted features (Bay et al. 2008; Lowe 2004), deep features (DeTone, Malisiewicz, and Rabinovich 2018; Dusmanu et al. 2019; Germain, Bourmaud, and Lepetit 2020; Sarlin et al. 2020) have been extensively utilized recently to improve feature matching for better localization. Very recently, some state-of-the-art (SOTA) feature matching methods have been proposed to train the deep features and align features through pose refinement in an end-to-end manner (Lindenberger et al. 2021; Sarlin et al. 2021). However, these structure-based methods rely on 2D-3D or 2D-2D point matching, and thus the accuracy is limited when the feature extraction and matching is sparse or noisy due to the large view changes between images and textureless structures.

Regression-based localization trains a neural network and takes the network parameters as a global map representation which can directly regress the 6-DOF camera poses (Kendall, Grimes, and Cipolla 2015; Balntas, Li, and Prisacariu 2018; Kendall and Cipolla 2017; Moreau et al. 2022a; Shavit, Ferens, and Keller 2021) or the 3D scene coordinate of each pixel (Cavallari et al. 2017; Li et al. 2020; Yang et al. 2019) by taking the query image as the network input. For the simplicity and end-to-end training manner, these methods have attracted considerable attention. However, these methods are usually scene-specific, and the accuracy heavily relies on the distribution of the training images with poor generalization to new viewpoints (Sarlin et al. 2021). Thus Neural Radiance Fields (NeRF) (Mildenhall et al. 2020) has been introduced recently to render realistic novel viewpoint images for data augmentation (Chen et al. 2022; Chen, Wang, and Prisacariu 2021; Moreau et al. 2022b) that can boost the training of the regression network. However, these methods are inherently regression-based, which imposes constraints on the localization accuracy as it is not feasible to indefinitely expand the training data and cover the whole 6D pose space.

To fix the problems of existing methods, we propose a novel framework for visual localization, *i.e.*, PNeRFLoc, that integrates the structure-based framework and the rendering-based optimization with NeRF representation. Specifically, on one hand, our framework supports the initial pose estimation by matching 2D and 3D feature points; on the other hand, it also enables the pose refinement with novel view synthesis using rendering-based optimization, *i.e.*, minimizing the

*Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

photometric error between the rendered image and the query image. In this way, compared to the existing NeRF-boosted methods (Chen et al. 2022; Chen, Wang, and Prisacariu 2021; Moreau et al. 2022b), our approach transcends the limitations of regression-based techniques, achieving significant accuracy improvements in both indoor and outdoor scenes. Moreover, compared to the SOTA feature matching methods which may be limited by sparse matches between reference and query images due to large view changes and thus stuck in local optima, our method can achieve better accuracy by minimizing the photometric loss with the capability to render novel-view images.

However, it is non-trivial to design the framework. First, there is no unified scene representation that supports both 2D-3D feature matching used in structure-based localization and neural rendering for rendering-based optimization. In this paper, we adapt a recent point-based neural radiance field representation (*i.e.*, PointNeRF (Xu et al. 2022)) and design a feature adaptation module to bridge the gap between the scene-agnostic features for localization and the scene-specific features for neural rendering. We find that although these two types of features aim for different tasks, they can be easily transferred via a feature adaptation module. In this way, we can utilize any existing scene-agnostic features for initial localization (*e.g.*, R2D2 (Revaud et al. 2019)), and learn the scene-specific adaptation module together with the NeRF models. Moreover, from the adaptation module, we can also learn a score for each dense feature for better feature matching and initial localization. Second, the rendering-based optimization may be easily stuck in the local minimum (Maggio et al. 2022) due to the backpropagation through the networks and also time-consuming. To improve the neural rendering-based optimization with point-based representation, we further propose a novel efficient rendering-based optimization framework by aligning the rendered image with the query image and minimizing the warping loss function. In this way, we don't need to render a new image for each step of optimization and avoid the backpropagation through the networks for better convergence. Lastly, to further improve the robustness of the proposed method for outdoor illumination changes and dynamic objects, we utilize appearance embedding and segmentation masks to handle varying lighting conditions and complex occlusions respectively.

Our contributions can be summarized as follows. At first, we propose a novel visual localization framework with a unified scene representation, *i.e.*, PNeRFLoc, which enables both structure-based estimation and render-based optimization for robust and accurate pose estimation. Second, to close the gaps between the features for visual localization and neural rendering, we propose a novel feature adaptation module that can be learned together with NeRF models. Furthermore, a novel efficient rendering-based framework with a warping loss function is proposed to improve the efficacy and efficiency of neural rendering-based optimization. Extensive experiments show that the proposed framework outperforms existing learning-based methods when the NeRF model can be well learned, and performs on-par with the SOTA method on the visual localization benchmark dataset.

Related Work

Structure-based localization. Structure-based methods (Camposco et al. 2017; Cheng et al. 2019; Sattler et al. 2015; Sattler, Leibe, and Kobbelt 2016; Toft et al. 2018; Zeisl, Sattler, and Pollefeys 2015) utilize 3D scene information from structure from motion (SfM), and a query image taken from the same scene can be registered with explicit 2D-3D correspondences and PnP + RANSAC algorithm. Typically, these methods can yield accurate poses but are prone to noisy matches. To mitigate outlier influence, recent scene coordinates regression (Brachmann et al. 2017; Brachmann and Rother 2021; Yang et al. 2019) methods rely on CNNs to fuse semantic features for obtaining accurate dense correspondences map, while the recent (Sarlin et al. 2020) excels graphical transformer with 2D relative positional encoding to achieve impressive sparse matching results. Despite the promising performance, the structure-based methods still suffer large-view changes especially when a few reference points are available.

Regression-based localization. PoseNet (Kendall, Grimes, and Cipolla 2015) and its subsequent work (Walch et al. 2017) regress the camera pose of an image directly through CNN or LSTM. These methods are limited in terms of scalability and performance. Despite some attempts to improve the accuracy by incorporating geometry prior (Brahmbhatt et al. 2018), these methods can only perform comparable results to that of image retrieval baselines (Arandjelovic et al. 2016; Torii et al. 2015) and cannot achieve the identical performance of structure-based counterparts. Moreover, adapting these regressed models to novel scenes is prohibited, which narrows their potential for real-time applications.

Localization with NeRF. Neural Radiance Fields (Mildenhall et al. 2020) has recently been employed for localization tasks. This is because NeRF can synthesize high-quality novel view images, which can be beneficial for localization tasks. For example, Purkait *et al.* proposed LENS (Moreau et al. 2022b), which uses NeRF-w (Martin-Brualla et al. 2021) to render realistic synthetic images to expand the training space. LENS leverages the NeRF-w model to obtain scene geometry information and render views from virtual camera poses covering the entire scene. However, LENS is limited by its long-time offline pre-training and infeasibility of covering the whole pose space, and it also lacks compensation for the domain gap between synthetic and real images, such as pedestrians and vehicles in outdoor scenes. Chen *et al.* proposed DFNet (Chen et al. 2022), which incorporates an additional feature extractor to learn high-level features to bridge the domain gap between synthetic and real images. However, the training process remains lengthy, as DFNet still needs to train NeRF, pose regression, and feature extraction networks separately. Maggio *et al.* proposed a Monte Carlo localization method called Loc-NeRF (Maggio et al. 2022), where Loc-NeRF continuously samples candidate poses under the initial pose and uses NeRF to render novel views to find the correct pose direction. However, Loc-NeRF is unstable and still requires an initial camera pose. Moreover, Yen-Chen *et al.* introduced iNeRF (Yen-Chen et al. 2021), an inverse NeRF approach to optimize camera poses, but it is also limited by the need to provide an initial pose.

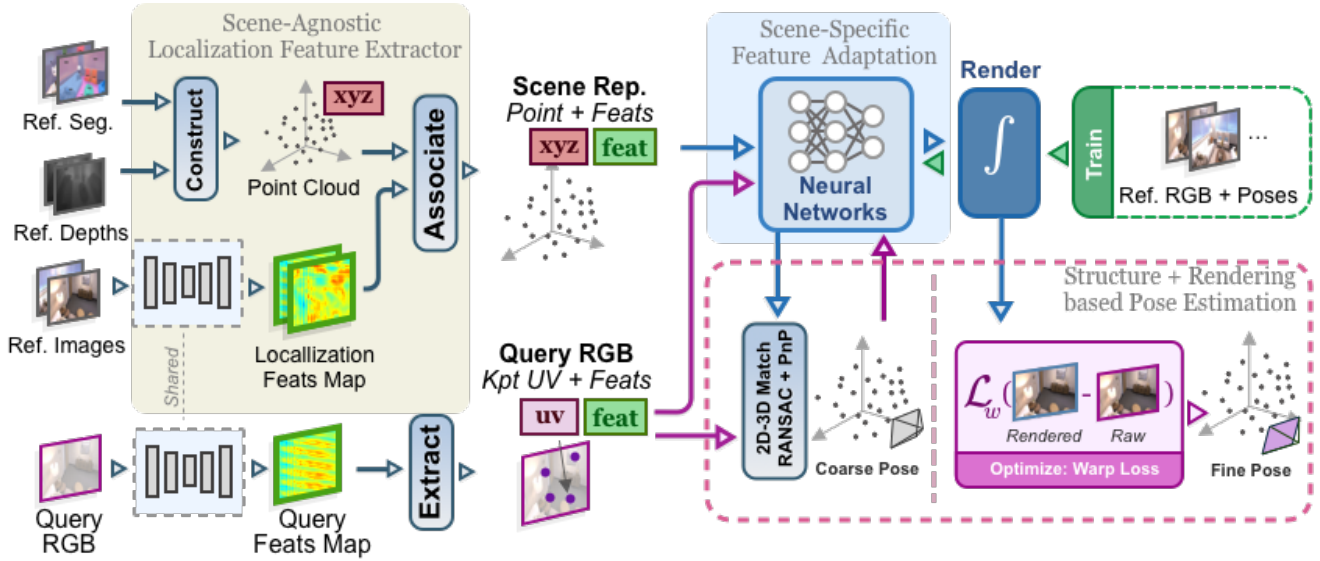


Figure 1: Visual localization with PNeRFLoc. In the proposed framework, we associate raw point clouds with scene-agnostic localization features and train a scene-specific feature adaptation together with the point-based neural radiance fields. Subsequently, PNeRFLoc integrates structure-based localization with novel rendering-based optimization to accurately estimate the 6-DOF camera pose of the query image.

Method

We propose a novel visual localization framework called **PNeRFLoc** based on the scene representation as shown in Fig.1. In order to enable both structure-based estimation and render-based optimization in a unified framework, we adapt the recent point-based radiance field representation (Xu et al. 2022) and design a feature adaptation module to bridge the scene-agnostic localization feature and the point-based neural rendering (Sec. 3.2). Additionally, to prevent iterative re-rendering of images for every optimization step like iNeRF (Yen-Chen et al. 2021), we propose an efficient rendering-based optimization strategy by minimizing the warping loss function to align the pixels on the rendered image and the query image, which reduced the neural rendering frequency to just once for most cases, while performing high accuracy (Sec. 3.3).

Point-based Radiance Field Representation

Neural Radiance Fields (NeRF) compute pixel radiance by sampling points along the ray shot through each pixel and computing the integral result. Specifically, each pixel in an image corresponds to a ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$. To render the color of ray \mathbf{r} , NeRF draws the point samples with distances $\{t_i\}_{i=1}^N$ to the camera origin \mathbf{o} along the ray, and passes the point locations $\mathbf{r}(t_i)$ as well as view directions \mathbf{d} to obtain density σ_i and colors \mathbf{c}_i . The resulting color is rendered following the quadrature rules (Max 1995):

$$\hat{\mathbf{C}}(\mathbf{r}) = \mathcal{R}(\mathbf{r}, \mathbf{c}, \sigma) = \sum_{k=1}^K T(t_k) \alpha(\sigma(t_k) \delta(t_k)) \mathbf{c}(t_k), \quad (1)$$

$$T(t_k) = \exp\left(-\sum_{k'=1}^{k-1} \sigma(t_{k'}) \delta_{k'}\right), \quad \alpha(x) = 1 - \exp(-x),$$

where $\mathcal{R}(\mathbf{r}, \mathbf{c}, \sigma)$ is the volumetric rendering through ray \mathbf{r} of color \mathbf{c} with density σ , $\mathbf{c}(t)$ and $\sigma(t)$ are the color and density at point $\mathbf{r}(t)$ respectively, and $\delta_k = t_{k+1} - t_k$ is the distance between two adjacent sampling points on the ray. Stratified sampling and informed sampling are used to select sample points $\{t_k\}_{k=1}^K$ between the near plane t_n and far plane t_f . Additionally, the depth \hat{D} of each ray \mathbf{r} can be computed as:

$$\hat{D} = \sum_{k=1}^K T(t_k) \alpha(\sigma(t_k) \delta(t_k)) t_k. \quad (2)$$

Following PointNeRF (Xu et al. 2022), we regress the radiance field from the point cloud $P = \{(p_i, f_i, \gamma_i) | i = 1, \dots, N\}$, where each point i is located at p_i and associated with a feature vector f_i that encodes the local scene content. And γ_i represents the confidence of a point being located on the actual surface of the scene. Given any 3D location x , we query K neighboring neural points around x and regress the density σ and view-dependent color \mathbf{c} from any viewing direction \mathbf{d} as:

$$(\sigma, \mathbf{c}) = \mathbf{PointNeRF}(x, \mathbf{d}, p_1, f_1, \gamma_1, \dots, p_K, f_K, \gamma_K). \quad (3)$$

To enable PointNeRF to handle dynamic objects and illumination changes, we adopt the appearance embedding from NeRF-W (Martin-Brualla et al. 2021) and a segmentation mask to handle occasional object occlusions and illumination variations. Contrary to NeRF-W which directly employs a transient MLP to address the issue of occasional object occlusions, we adopt a more stable approach by utilizing a segmentation mask to compel the network to focus exclusively on architectural areas. In this paper, we utilize Detectron2¹ to perform object detection on the images and

¹<https://github.com/facebookresearch/detectron2>

generate segmentation masks.

Scene-Specific Feature Adaptation

Once the point-based NeRF model is built, a straightforward way is to utilize the learned point-wise neural features for feature matching. However, we find that these neural features are not distinctive enough and cannot be used as robust and efficient descriptors for feature matching as shown in our supplementary material because these features are learned to encode local color and geometry information of the specific scene for the neural rendering. Based on this observation, we resort to existing well-studied deep features for visual localization that are trained on large datasets for feature matching and design a feature adaptation module to bridge the features for visual localization and neural rendering. Moreover, we can also learn the scores of neural features with the adaptation module for better feature matching for the specific scene.

Scene-agnostic point localization feature extractor. In this paper, we utilize the deep feature R2D2 (Revaud et al. 2019) as the scene-agnostic feature for visual localization due to its ability to robustly extract reliable and distinctive features. For each reference image $I_k \in \mathbb{R}^{W \times H \times 3}$, the R2D2 network extracts a feature map $\mathbf{F}_k \in \mathbb{R}^{W \times H \times 128}$. For each 3D point i constructed from each reference image k , we define the scene-agnostic point feature as:

$$f_i = \mathbf{F}_k[p_i] \in \mathbb{R}^{128}, \quad (4)$$

where p_i is the projection of i in the reference image and $[\cdot]$ is a lookup with sub-pixel interpolation. Searching for matches throughout the entire point cloud is inefficient as the reliability of scene-agnostic localization features can be compromised by the structure of the scene. Therefore, we utilize the matching score of each point (introduced in the feature adaptation) to enable point filtering during the matching process. By removing candidate points in the point cloud below a certain score threshold, we can reduce the number of matching pairs to be computed, thus improving efficiency.

Scene-specific feature adaptation. As explained before, due to the significant gap between scene-agnostic features for localization and scene-specific features for neural rendering, we cannot learn the radiance fields from the R2D2 features. Thus we design a feature adaptation module to bridge this gap, which consists of a four-layer Multi-Layer Perceptron (MLP). We empirically find that despite the scene-agnostic feature and the scene-specific NeRF representation feature aiming at two completely different tasks, they can be adapted via the designed module. Thus any other SOTA scene-agnostic feature for visual localization can also be utilized in our framework. Moreover, as mentioned above, we also utilize the adaptation module to learn a score S for each point in the point cloud according to its dense features and position. These scores are then used for point filtering, which improves the efficiency of feature matching while maintaining the accuracy of the final pose estimation.

Scene-specific PointNeRF reconstruction. Given the pre-trained point localization feature extractor and feature adaptation module, similar to PointNeRF (Xu et al. 2022), we learn

the NeRF model by minimizing the following loss function:

$$\mathcal{L}_{render} = \sum_{r \in R} \|\hat{C}(r) - C(r)\|_2^2, \quad (5)$$

where R is the set of rays in each batch, and $C(r)$, $\hat{C}(r)$ are the ground truth and predicted RGB colors for ray r computed by Eq.1. To be noted, we also learned to fine-tune the feature adaptation module for each scene for better rendering quality. Please refer to our supp. material for more details.

Two-stage Pose Estimation

Once the NeRF model is learned, we design a two-stage pose estimation framework for the query image during the test.

Initialization with structure-based localization. The goal of the structure-based localization stage is to establish the correspondence between the 2D key points on the query image and the 3D points in the scene point cloud, thereby providing an initial pose estimate for the subsequent pose refinement stage. For each query image q with the keypoints P_q and features \mathbf{F}_q extracted by the scene-agnostic localization feature extractor, and the point cloud P_r generated by PointNeRF with features \mathbf{F}_r , we can find a 2d-3D correspondence:

$$\forall i \in P_q, \quad M(i) = \arg \max_{j \in P_r} \frac{\mathbf{F}_q^i \cdot \mathbf{F}_r^j}{\|\mathbf{F}_q^i\| \|\mathbf{F}_r^j\|}, \quad (6)$$

where $M(i)$ signifies the corresponding point within the point cloud for the keypoint i present on the query image q , which is ascertained via the maximization of cosine similarity. However, as mentioned in Sec. 3.2, the process of directly seeking correspondences within the entire point cloud proves to be inefficient. In response to this, we employ a thresholding technique based on the learned score S to filter the point cloud P_r . Given a threshold S_t , we can get the filtered point cloud as $P_s = \{i \in P_r \mid S_i \geq S_t\}$, where S_i denotes the learned score for the point i . For each query image q and the discovered correspondence M , we define a residual:

$$r_i = \|p_i - \prod(\mathbf{R}M(i) + \mathbf{t})\|_2, \quad (7)$$

where $\prod(\cdot)$ represents the pixel obtained post the projection of the 3D point onto the image. (\mathbf{R}, \mathbf{t}) denote the camera pose to be determined, while p_i signifies the pixel of the keypoint i within the query image. The total error over all key points is:

$$E(\mathbf{R}, \mathbf{t}) = \sum_{i \in P_q} r_i. \quad (8)$$

Moreover, a direct optimization of Eq.8 is often susceptible to the distortions caused by incorrect correspondences (outliers). Therefore, a RANSAC loop is also adopted, effectively improving the accuracy.

Pose refinement with efficient rendering-based optimization. Previous works (Yen-Chen et al. 2021; Zhu et al. 2022) have utilized gradient descent to minimize the photometric residuals between the rendered and input images for local pose estimation. However, this optimization method is inefficient since neural rendering is required for each optimization step, and it is also unstable due to the backpropagation over the deep networks. Therefore, we propose a novel and efficient rendering-based optimization strategy using the warping

loss function, which only requires rendering the image once and avoiding the backpropagation through the networks.

Specifically, for a given query image q and initial pose (\mathbf{R}, \mathbf{t}) , PNeRFLoc first renders the visual reference image q_r under the initial pose according to Eq. (1) and the depth map d_q according to Eq. (2). Subsequently, we randomly sample N pixels within the image q_r . For the pose $(\mathbf{R}', \mathbf{t}')$ that we aspire to optimize, we define the warping loss function as:

$$\mathcal{L}_{warping} = \sum_{p_i \in N} \|C(q, W(p_i, \mathbf{R}, \mathbf{t}, \mathbf{R}', \mathbf{t}')) - C(q_r, p_i)\|_2, \quad (9)$$

$$W(p_i, \mathbf{R}, \mathbf{t}, \mathbf{R}', \mathbf{t}') = \prod(\mathbf{R}'(\mathbf{R}^{-1} \prod^{-1}(p_i, \hat{D}(p_i) - \mathbf{R}^{-1}\mathbf{t}) + \mathbf{t}')), \quad (10)$$

where $C(q_r, p_i)$ represents the RGB color at pixel p_i on rendering image q_r , and the function W denotes the corresponding pixel on query image q by warping p_i from render image q_r . Specifically, W back-projects p_i into the 3D space of the q_r 's camera coordinate system using the depth $\hat{D}(p_i)$, and then transposes it to the camera coordinate system of image q through the camera pose $(\mathbf{R}', \mathbf{t}')$ and projects it onto image q finally. However, we find that there are often blanks in the rendering images, which occur when the rays emitted from the camera pass through the gaps in the point cloud and do not aggregate to the neural points. In this case, incorrect depth and color can interfere with the optimization. Hence, we propose using a blank depth mask to handle such situations. For the set of sampled pixels N on the visual reference image q_r , we define the valid pixels set as $N_v = \{p_i \in N | \hat{D}(p_i) \geq 0.01\}$ and let the warping loss function only consider the pixels in N_v . Our rendering-based optimization method optimizes the pose $(\mathbf{R}', \mathbf{t}')$ by using warping loss aligning the RGB colors of sampled pixels on q_r and q . Thereby we avoid gradient descent through the complex neural networks and improve accuracy and efficiency. Moreover, when the viewpoint changes significantly between the visual reference image and the query image, the optimization result may not reach its optimum. We could potentially enhance the accuracy by iteratively rendering the visual reference image multiple times. However, we found that a single rendering's outcome was already satisfactory in our experiments. Therefore, to save time, our optimization process only renders the visual reference image once in practice.

Experiments

We first compare our method with various representative and SOTA learning approaches (Sarlin et al. 2021; Moreau et al. 2022a,b; Brachmann and Rother 2021) on both synthetic datasets and real-world datasets. Then, we offer insights into PNeRFLoc through additional ablation experiments.

Datasets and Implementation Details

Datasets. Following (Chen et al. 2022; Moreau et al. 2022b,a), we evaluate our method on two standard localization datasets since they have well-distributed training images which support dense 3D reconstruction. Moreover, we generate a synthetic localization dataset using the commonly used Replica dataset in NeRF-based SLAM systems.

	Methods	CoordiNet	PixLoc	PNeRFLoc
Replica	room0	1.60 / 50.8	0.055 / 1.89	0.005 / 0.29
	room1	1.38 / 47.3	0.020 / 0.36	0.016 / 0.55
	room2	1.26 / 20.2	0.901 / 8.71	0.022 / 0.92
	office0	1.14 / 20.1	0.021 / 0.71	0.006 / 0.69
	office1	0.81 / 36.3	0.016 / 0.75	0.017 / 0.64
	office2	0.83 / 19.9	0.012 / 0.40	0.007 / 0.44
	office3	0.76 / 18.8	0.015 / 0.67	0.006 / 0.30
	office4	0.89 / 46.3	0.033 / 0.82	0.009 / 0.23

Table 1: Comparison on Replica datasets. We report median translation/rotation errors (meters/degrees) and the best results are highlighted in bold.

- **Cambridge Landmarks (Kendall, Grimes, and Cipolla 2015)** contains five outdoor scenes, with 200 to 2000 images captured at different times for each scene. This dataset is challenging for camera pose estimation because the query images are taken at different times than the reference images, resulting in different lighting conditions and occlusions from objects such as people and vehicles.
- **7Scenes (Shotton et al. 2013)** contains seven indoor scenes, captured by a Kinect RGB-D sensor. Each scene has 1k to 7k reference images and 1k to 5k query images, captured along different trajectories.
- **Replica (Straub et al. 2019)** contains eight synthetic indoor scenes, commonly used for SLAM evaluation. We follow iMAP (Sucar et al. 2021), using its produced sequences as training images, with an image size of 1200*680 pixels, and randomly generate 50-120 query images. Due to the small number of reference images and the significant changes in the viewpoint of the query images, this dataset presents a certain level of challenge for localization tasks.

Implementation. We use R2D2 (Revaud et al. 2019) as the scene-agnostic localization feature extractor. In the structure-based localization stage, the score threshold S_t is set to 0.7, and the number of RANSAC iterations is set to 20k. During the rendering-based localization stage, we use the Adam optimizer with a learning rate of 0.001. All our experiments are evaluated on a single NVIDIA GeForce RTX 3090 GPU. Similar to DSAC* (Brachmann and Rother 2021), we obtain the estimated depth images rendered from a 3D model learned by Factorized-NeRF (Zhao et al. 2022). For the 7scenes dataset, we follow PixLoc (Sarlin et al. 2021) utilizing the estimated depth rendered by DSAC* (Brachmann and Rother 2021). Lastly, we leveraged a pre-trained model from MonoSDF (Yu et al. 2022) to render the estimated depth for the Replica dataset. Please refer to the supp. material for more details.

Evaluation on the Replica Dataset

We first compare our method with the SOTA structure-based method PixLoc (Sarlin et al. 2021) and the regression-based method CoordiNet (Moreau et al. 2022a) on the Replica dataset. Our method and PixLoc adopt 200 images as reference images, while for CoordiNet, we use 2000 images for training to generate reasonable results. The evaluation results




Figure 2: Pose estimation results. We visualize the rendering images based on the estimated pose at time t and the query image to compare different optimization methods.

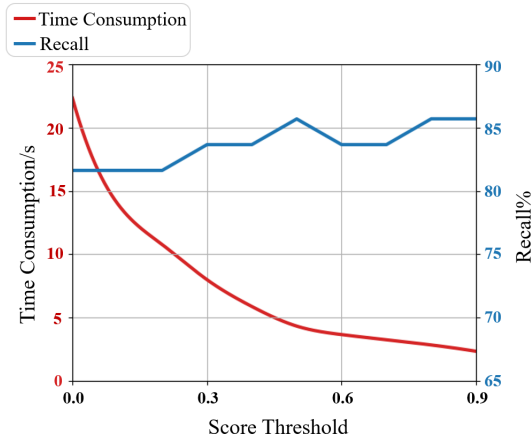


Figure 3: We show the effectiveness of our reliable and repeatable score filtering on localization efficiency and accuracy.

are shown in Table 1. Our method achieves state-of-the-art results on the Replica dataset. We believe that PNeRFLoc surpasses PixLoc on the Replica dataset for two primary reasons: i) Each query image in Replica has a relatively large view-point change compared to the reference image, leading to a large initial reprojection error in PixLoc and causing the optimization to fall into incorrect local minima; ii) with high-quality input images and accurate camera poses, PNeRFLoc can learn a fine NeRF model, which further facilitates rendering-based optimization given the initial pose estimation from the structure-based localization. The regression-based method CoordiNet has the worst performance unsurprisingly due to its poor generalization to the query image with large view-point changes although more reference images are provided to train the regression model. Please refer to our supplementary material for more comparisons.

Evaluation on the Cambridge and 7Scenes

We compare with multiple SOTA approaches (Sarlin et al. 2021; Moreau et al. 2022a,b; Brachmann and Rother 2021) on the benchmark visual localization datasets, *i.e.*, Cambridge Landmarks and 7Scenes. We report the median translation

and rotation error for each scene in Table 2.

For the indoor 7Scenes dataset, since the generated depth by DSAC* (Brachmann and Rother 2021) are noisy with misalignments, it affects the training of the methods based on depth images, like DSAC* and our method. As a result, our method performs on-par or slightly worse than the SOTA PixLoc method on 7Scenes, while it still outperforms all other methods in general. In the subsequent ablation experiments, we provide the results with relatively better depth inputs, which confirms our observations.

For the outdoor Cambridge Landmarks dataset, there are large appearance variations and dynamic objects. Moreover, we find that the provided camera poses and intrinsic parameters of the training images are not very accurate, which prevents PNeRFLoc from learning a fine 3D NeRF model and rendering higher-quality novel view images. Even so, PNeRFLoc still performs on-par with the SOTA method and performs much better than all other NeRF-boosted localization methods (Moreau et al. 2022b; Chen et al. 2022), which demonstrates its robustness and the potential of the NeRF-based methods for outdoor datasets. As long as more accurate depths and camera poses are provided, our method achieves the SOTA performance as shown on the Replica dataset.

Ablation Studies

Justification of the proposed rendering-based optimization. As shown in Fig. 2 and Table 3, we justify our design decisions by comparing different variants of PNeRFLoc. All experiments were optimized 250 times during the rendering-based localization stage. We report the median translation/rotation errors (meters/degrees) and time consumption (seconds/per image). We can see that the proposed rendering-based optimization significantly improves the localization accuracy given the initial pose estimated in the structure-based estimation stage. Without the blank depth mask, the accuracy slightly degrades due to the numerous sampling points on the image, and a small proportion of blank area sampling does not affect the overall trend of optimization. Furthermore, direct optimization using photometric loss is more time-consuming, and it may also fall into incorrect local minima due to the backpropagation through networks. These ablation studies demonstrate the efficacy and efficiency of

	Methods	PoseNet	CoordiNet	LENS	DFNet	DSAC*	PixLoc	PNeRF Loc
7Scenes	Chess	0.32 / 8.12	0.14 / 6.7	0.03 / 1.3	0.04 / 1.48	0.02 / 1.10	0.02 / 0.80	0.02 / 0.80
	Fire	0.47 / 14.4	0.27 / 11.6	0.10 / 3.7	0.04 / 2.16	0.02 / 1.24	0.02 / 0.73	0.02 / 0.88
	Heads	0.29 / 12.0	0.13 / 13.6	0.07 / 5.8	0.03 / 1.82	0.01 / 1.82	0.01 / 0.82	0.01 / 0.83
	Office	0.48 / 7.68	0.21 / 8.6	0.07 / 1.9	0.07 / 2.01	0.03 / 1.15	0.03 / 0.82	0.03 / 1.05
	Pumpkin	0.47 / 8.42	0.25 / 7.2	0.08 / 2.2	0.09 / 2.26	0.04 / 1.34	0.04 / 1.21	0.06 / 1.51
	Kitchen	0.59 / 8.64	0.26 / 7.5	0.09 / 2.2	0.09 / 2.42	0.04 / 1.68	0.03 / 1.20	0.05 / 1.54
	Stairs	0.47 / 13.8	0.28 / 12.9	0.14 / 3.6	0.14 / 3.31	0.03 / 1.16	0.05 / 1.30	0.32 / 5.73
Cambridge	Kings	1.66 / 4.86	0.70 / 2.92	0.33 / 0.5	0.43 / 0.87	0.15 / 0.3	0.14 / 0.24	0.24 / 0.29
	Hospital	2.62 / 4.90	0.97 / 2.08	0.44 / 0.9	0.46 / 0.87	0.21 / 0.4	0.16 / 0.32	0.28 / 0.37
	Shop	1.41 / 7.18	0.73 / 4.69	0.27 / 1.6	0.16 / 0.59	0.05 / 0.3	0.05 / 0.23	0.06 / 0.27
	Church	2.45 / 7.96	1.32 / 3.56	0.53 / 1.6	0.50 / 1.49	0.13 / 0.4	0.10 / 0.34	0.40 / 0.55
	Court	2.45 / 3.98	-	-	-	0.49 / 0.3	0.30 / 0.14	0.81 / 0.25

Table 2: Comparison on the Cambridge Landmarks and 7Scenes datasets. We report median translation/rotation errors(meters/degrees). We highlight the top two results in bold.

Config.	Replica	
	room0	office0
w/o Rendering-Based Optimization	0.030 / 0.79 / 3.20	0.082 / 1.38 / 1.88
w/o Blank Depth Mask	0.027 / 0.63 / 5.84	0.050 / 1.09 / 5.66
w/o Warping Loss, w/ Photometric Loss	0.035 / 0.81 / 47.7	0.082 / 1.38 / 39.2
Full Model	0.005 / 0.29 / 5.56	0.006 / 0.56 / 5.45

Table 3: Ablation study. We report the median translation/rotation errors (meters/degrees) and time consumption (seconds/per image). The best results are highlighted in bold.

Config.	7Scenes			Replica		
	chess	office	stairs	room1	office0	office1
With Estimated Depth	0.02 / 0.80	0.03 / 1.05	0.32 / 5.73	0.02 / 0.55	0.01 / 0.56	0.02 / 0.64
With GT Depth	0.02 / 0.80	0.03 / 1.06	0.20 / 3.61	0.01 / 0.34	0.01 / 0.54	0.01 / 0.39

Table 4: Comparison of using input depth and estimated depth. We report median translation/rotation errors (meters/degrees) and the best results are highlighted in bold.

the proposed rendering-based optimization.

Impact of score filtering. We analyzed the impact of score filtering on the Replica’s office4 scene. As shown in Fig.3, we report the time consumption (in seconds) and recall at (5cm, 5°). With the increase of score threshold, the efficiency of the PnP algorithm is significantly improved, which is because the number of remaining candidate points in the point cloud decreases, saving the time to calculate cosine similarity. At the same time, we can see that the accuracy is preserved and even slightly improved with the score filtering because more reliable points for the scene are selected.

Performance with input depth images. Since our method requires depth images to establish a point-based NeRF representation of the scene, we analyze the robustness of PixelLoc against the depth images. So as shown in Table ??, we compare our results of learning the NeRF model with ground-truth input depth and estimated depth on the 7Scenes and Replica datasets. Since the Replica dataset is a synthetic dataset, its GT depth is dense and accurate, allowing for more precise neural point clouds and better rendering quality. In

contrast, the estimated depth often loses details. As illustrated in Fig.4, the estimated depth rendered by MonoSDF (Yu et al. 2022) fails to capture the vase. Therefore, using GT depth on the Replica dataset significantly improves localization accuracy. In real-world scenes, however, the depth obtained by the Kinect RGB-D sensor is noisy, which is mainly influenced by the reflection and refraction of object surfaces, as well as the sensor’s maximum and minimum measurement ranges. Consequently, the improvement in scene training quality is limited, and there is no significant improvement in localization accuracy. However, the advantage of taking the input depth is evident in the stairs scene, where the complex spatial structure leads to misalignment in the estimated depth. Please refer to our supp. material for more ablation studies.

Conclusion

In this paper, we present a novel visual localization method based on point-based neural scene representation. With a novel feature adaption module that bridges the features for

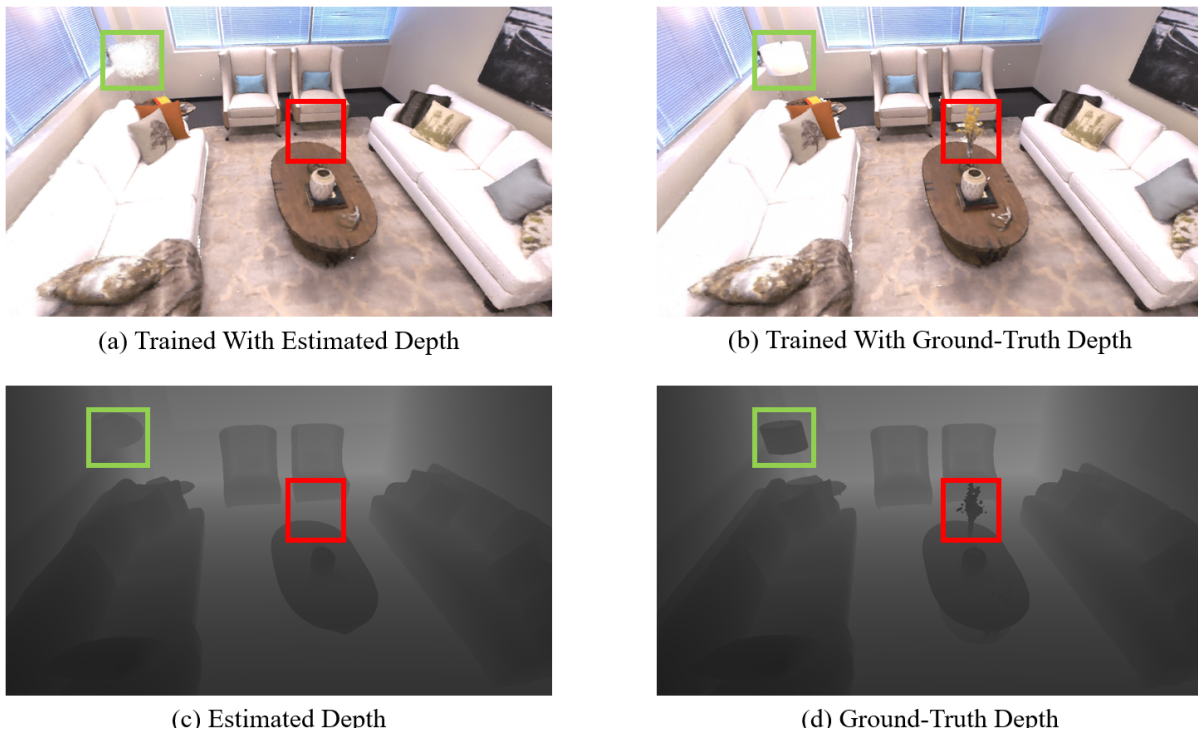


Figure 4: We inspect the trained model with estimated depth or with input GT depth.

localization and neural rendering, the proposed PNeRFLoc enables 2D-3D feature matching for initial pose estimation and rendering-based optimization for pose refinement. Moreover, we also develop several techniques for efficient rendering-based optimization and robustness against illumination changes and dynamic objects. Experiments show the superiority of the proposed method by integrating both structure-based and rendering-based optimization, especially on the synthetic data suitable for NeRF modeling. Although our current framework is more efficient than the existing neural rendering-based optimization, we should further improve the efficiency and integrate it into visual odometry for real-time applications.

Acknowledgments

This work was partially supported by the NSFC (No. 62102356). We are also very grateful for the illustrations crafted by Lin Zeng.

References

- Arandjelovic, R.; Gronat, P.; Torii, A.; Pajdla, T.; and Sivic, J. 2016. NetVLAD: CNN architecture for weakly supervised place recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5297–5307.
- Balntas, V.; Li, S.; and Prisacariu, V. 2018. Relocnet: Continuous metric learning relocalisation using neural nets. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 751–767.
- Bay, H.; Ess, A.; Tuytelaars, T.; and Van Gool, L. 2008. Speeded-up robust features (SURF). *Computer vision and image understanding*, 110(3): 346–359.
- Brachmann, E.; Krull, A.; Nowozin, S.; Shotton, J.; Michel, F.; Gumhold, S.; and Rother, C. 2017. Dsac-differentiable ransac for camera localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6684–6692.
- Brachmann, E.; and Rother, C. 2021. Visual camera relocalization from RGB and RGB-D images using DSAC. *IEEE transactions on pattern analysis and machine intelligence*, 44(9): 5847–5865.
- Brahmbhatt, S.; Gu, J.; Kim, K.; Hays, J.; and Kautz, J. 2018. Geometry-aware learning of maps for camera localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2616–2625.
- Bujnak, M.; Kukulova, Z.; and Pajdla, T. 2008. A general solution to the P4P problem for camera with unknown focal length. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, 1–8. IEEE.
- Camposeco, F.; Sattler, T.; Cohen, A.; Geiger, A.; and Pollefeys, M. 2017. Toroidal constraints for two-point localization under high outlier ratios. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4545–4553.
- Cavallari, T.; Golodetz, S.; Lord, N. A.; Valentin, J.; Di Stefano, L.; and Torr, P. H. 2017. On-the-fly adaptation of regression forests for online camera relocalisation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4457–4466.

- Chen, S.; Li, X.; Wang, Z.; and Prisacariu, V. A. 2022. Dfnet: Enhance absolute pose regression with direct feature matching. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part X*, 1–17. Springer.
- Chen, S.; Wang, Z.; and Prisacariu, V. 2021. Direct-PoseNet: absolute pose regression with photometric consistency. In *2021 International Conference on 3D Vision (3DV)*, 1175–1185. IEEE.
- Cheng, W.; Lin, W.; Chen, K.; and Zhang, X. 2019. Cascaded Parallel Filtering for Memory-Efficient Image-Based Localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Chum, O.; and Matas, J. 2008. Optimal randomized RANSAC. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(8): 1472–1482.
- DeTone, D.; Malisiewicz, T.; and Rabinovich, A. 2018. Superpoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 224–236.
- Dusmanu, M.; Rocco, I.; Pajdla, T.; Pollefeys, M.; Sivic, J.; Torii, A.; and Sattler, T. 2019. D2-net: A trainable cnn for joint description and detection of local features. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8092–8101.
- Fischler, M. A.; and Bolles, R. C. 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6): 381–395.
- Germain, H.; Bourmaud, G.; and Lepetit, V. 2020. S2dnet: Learning accurate correspondences for sparse-to-dense feature matching. *arXiv preprint arXiv:2004.01673*.
- Haralick, B. M.; Lee, C.-N.; Ottenberg, K.; and Nölle, M. 1994. Review and analysis of solutions of the three point perspective pose estimation problem. *International journal of computer vision*, 13: 331–356.
- Kendall, A.; and Cipolla, R. 2017. Geometric loss functions for camera pose regression with deep learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5974–5983.
- Kendall, A.; Grimes, M.; and Cipolla, R. 2015. Posenet: A convolutional network for real-time 6-dof camera relocalization. In *Proceedings of the IEEE international conference on computer vision*, 2938–2946.
- Li, X.; Wang, S.; Zhao, Y.; Verbeek, J.; and Kannala, J. 2020. Hierarchical scene coordinate classification and regression for visual localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11983–11992.
- Lindenberger, P.; Sarlin, P.-E.; Larsson, V.; and Pollefeys, M. 2021. Pixel-perfect structure-from-motion with featuremetric refinement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5987–5997.
- Lowe, D. G. 2004. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60: 91–110.
- Maggio, D.; Abate, M.; Shi, J.; Mario, C.; and Carlone, L. 2022. Loc-NeRF: Monte Carlo Localization using Neural Radiance Fields. *arXiv preprint arXiv:2209.09050*.
- Martin-Brualla, R.; Radwan, N.; Sajjadi, M. S.; Barron, J. T.; Dosovitskiy, A.; and Duckworth, D. 2021. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7210–7219.
- Max, N. 1995. Optical Models for Direct Volume Rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2): 99–108.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*.
- Moreau, A.; Piasco, N.; Tsishkou, D.; Stanculescu, B.; and de La Fortelle, A. 2022a. CoordiNet: uncertainty-aware pose regressor for reliable vehicle localization. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2229–2238.
- Moreau, A.; Piasco, N.; Tsishkou, D.; Stanculescu, B.; and de La Fortelle, A. 2022b. LENS: Localization enhanced by NeRF synthesis. In *Conference on Robot Learning*, 1347–1356. PMLR.
- Revaud, J.; Weinzaepfel, P.; de Souza, C. R.; and Humenberger, M. 2019. R2D2: Repeatable and Reliable Detector and Descriptor. In *NeurIPS*.
- Sarlin, P.-E.; Cadena, C.; Siegwart, R.; and Dymczyk, M. 2019. From Coarse to Fine: Robust Hierarchical Localization at Large Scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Sarlin, P.-E.; DeTone, D.; Malisiewicz, T.; and Rabinovich, A. 2020. SuperGlue: Learning feature matching with graph neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4938–4947.
- Sarlin, P.-E.; Unagar, A.; Larsson, M.; Germain, H.; Toft, C.; Larsson, V.; Pollefeys, M.; Lepetit, V.; Hammarstrand, L.; Kahl, F.; et al. 2021. Back to the feature: Learning robust camera localization from pixels to pose. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3247–3257.
- Sattler, T.; Havlena, M.; Radenovic, F.; Schindler, K.; and Pollefeys, M. 2015. Hyperpoints and fine vocabularies for large-scale location recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, 2102–2110.
- Sattler, T.; Leibe, B.; and Kobbelt, L. 2016. Efficient & effective prioritized matching for large-scale image-based localization. *IEEE transactions on pattern analysis and machine intelligence*, 39(9): 1744–1756.
- Shavit, Y.; Ferens, R.; and Keller, Y. 2021. Paying attention to activation maps in camera pose regression. *arXiv preprint arXiv:2103.11477*.
- Shotton, J.; Glocker, B.; Zach, C.; Izadi, S.; Criminisi, A.; and Fitzgibbon, A. 2013. Scene coordinate regression forests for camera relocalization in RGB-D images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2930–2937.

Straub, J.; Whelan, T.; Ma, L.; Chen, Y.; Wijmans, E.; Green, S.; Engel, J. J.; Mur-Artal, R.; Ren, C.; Verma, S.; et al. 2019. The Replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*.

Sucar, E.; Liu, S.; Ortiz, J.; and Davison, A. J. 2021. iMAP: Implicit mapping and positioning in real-time. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6229–6238.

Toft, C.; Stenborg, E.; Hammarstrand, L.; Brynte, L.; Pollefeys, M.; Sattler, T.; and Kahl, F. 2018. Semantic Match Consistency for Long-Term Visual Localization. In *Proceedings of the European Conference on Computer Vision (ECCV)*.

Torii, A.; Arandjelovic, R.; Sivic, J.; Okutomi, M.; and Pajdla, T. 2015. 24/7 place recognition by view synthesis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1808–1817.

Walch, F.; Hazirbas, C.; Leal-Taixe, L.; Sattler, T.; Hilsenbeck, S.; and Cremers, D. 2017. Image-based localization using lstms for structured feature correlation. In *Proceedings of the IEEE International Conference on Computer Vision*, 627–637.

Xu, Q.; Xu, Z.; Philip, J.; Bi, S.; Shu, Z.; Sunkavalli, K.; and Neumann, U. 2022. Point-nerf: Point-based neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5438–5448.

Yang, L.; Bai, Z.; Tang, C.; Li, H.; Furukawa, Y.; and Tan, P. 2019. Sanet: Scene agnostic network for camera localization. In *Proceedings of the IEEE/CVF international conference on computer vision*, 42–51.

Yen-Chen, L.; Florence, P.; Barron, J. T.; Rodriguez, A.; Isola, P.; and Lin, T.-Y. 2021. inerf: Inverting neural radiance fields for pose estimation. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1323–1330. IEEE.

Yu, Z.; Peng, S.; Niemeyer, M.; Sattler, T.; and Geiger, A. 2022. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *arXiv preprint arXiv:2206.00665*.

Zeisl, B.; Sattler, T.; and Pollefeys, M. 2015. Camera Pose Voting for Large-Scale Image-Based Localization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.

Zhao, B.; Yang, B.; Li, Z.; Li, Z.; Zhang, G.; Zhao, J.; Yin, D.; Cui, Z.; and Bao, H. 2022. Factorized and controllable neural re-rendering of outdoor scene for photo extrapolation. In *Proceedings of the 30th ACM International Conference on Multimedia*, 1455–1464.

Zhu, Z.; Peng, S.; Larsson, V.; Xu, W.; Bao, H.; Cui, Z.; Oswald, M. R.; and Pollefeys, M. 2022. Nice-slam: Neural implicit scalable encoding for slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12786–12796.