

# The NeRFect Match: Exploring NeRF Features for Visual Localization

Qunjie Zhou<sup>1\*</sup>, Maxim Maximov<sup>2\*</sup>, Or Litany<sup>13</sup>, and Laura Leal-Taixé<sup>1</sup>

<sup>1</sup> NVIDIA

<sup>2</sup> Technical University of Munich, Germany

<sup>3</sup> Technion, Israel

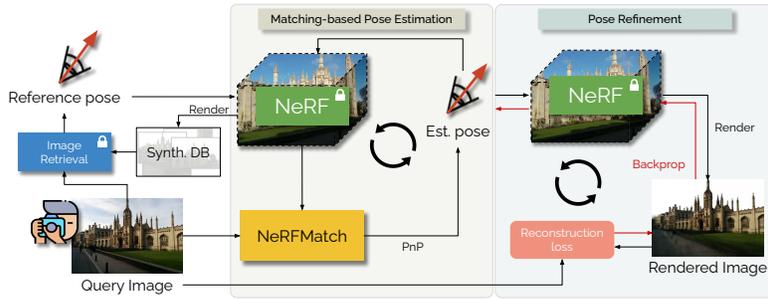
**Abstract.** In this work, we propose the use of Neural Radiance Fields (NeRF) as a scene representation for visual localization. Recently, NeRF has been employed to enhance pose regression and scene coordinate regression models by augmenting the training database, providing auxiliary supervision through rendered images, or serving as an iterative refinement module. We extend its recognized advantages – its ability to provide a compact scene representation with realistic appearances and accurate geometry – by exploring the potential of NeRF’s internal features in establishing precise 2D-3D matches for localization. To this end, we conduct a comprehensive examination of NeRF’s implicit knowledge, acquired through view synthesis, for matching under various conditions. This includes exploring different matching network architectures, extracting encoder features at multiple layers, and varying training configurations. Significantly, we introduce NeRFMatch, an advanced 2D-3D matching function that capitalizes on the internal knowledge of NeRF learned via view synthesis. Our evaluation of NeRFMatch on standard localization benchmarks, within a structure-based pipeline, achieves competitive results for localization performance on Cambridge Landmarks. We will release all models and code.

## 1 Introduction

Visual localization is the task of determining the camera pose of a query image w.r.t a 3D environment. Such ability to localize an agent in 3D is fundamental to applications such as robot navigation [20, 64], autonomous driving [26], and augmented reality [2, 61]. Different localization solutions can be categorized based on their underlying scene representation. Image retrieval [1, 4, 24, 60] represents a scene as a database of reference images with known camera poses. A more compact representation utilizes a 3D point cloud, where each point is triangulated from its 2D projections in multiple views. Structure-based methods [28, 36, 48, 51, 57] rely on such a 3D model with associated keypoint descriptors to perform accurate localization. Recently, MeshLoc [44] expanded structure-based localization by integrating dense 3D meshes, allowing to switch between different descriptors to establish the 2D-3D correspondences.

---

\* These authors contributed equally to this work



**Fig. 1: NeRF-based localization overview.** In this work, we propose to use NeRF as our scene representation for visual localization. Given a query image, we first retrieve its nearest reference pose using image retrieval, then use NeRFMatch to establish 2D-3D correspondences between the query image and the NeRF scene points to compute an initial pose estimate and finally improve its accuracy via pose refinement.

Different from the aforementioned *explicit* scene representations, absolute pose regression (APR) [5, 11, 30, 31, 53, 62] and scene coordinate regression (SCR) methods [6, 8–10, 35, 54, 67] learn to encode scene information within network parameters, either directly in the localization network, or as a separate collection of latent codes as learned scene representation [59]. Despite being more compact than a sparse point cloud and a 3D mesh, learned *implicit* scene representations are less interpretable and are limited to the task of visual localization.

Recently, Neural Radiance Fields (NeRF) [40] have emerged as a powerful representation of 3D scenes, encoded as continuous mapping of spatial coordinates and viewing angle to density and radiance. NeRF present several benefits: high interpretability, *i.e.*, one can easily render scene appearance and depth at any given viewpoint, as well as being highly compact, *e.g.*, a Mip-NeRF [3] model of 5.28 MB can represent a scene with a spatial extent ranging from  $1m^2$  [54] to  $5km^2$  [31] (Sec. 5). Owing to its attractive properties, NeRF is emerging as a *prime* 3D scene representation [65], alongside meshes, point clouds, and multi-view images, and has been applied to various other computer vision tasks such as semantic segmentation [22, 34, 71], 3D object detection [27, 66], Simultaneous Localization and Mapping (SLAM) [46, 55, 70, 74] and vision-based localization [15–18, 37, 39, 41, 42, 68].

NeRF has been leveraged for tackling visual localization in various ways. iNeRF-style approaches [16, 68] utilize a pre-trained NeRF as an inference-time pose refinement. Yet such methods commonly suffer from slow convergence and require pose initialization to be provided. The end-to-end APR and SCR methods use a pre-trained NeRF only at train-time to augment training samples [15, 42], provide consistency supervision [17, 18], and generate proxy depth ground-truth [15]. In this case, NeRF merely serves as an auxiliary representation. Orthogonal to the above ways of leveraging NeRF, recent works [37, 41] propose to use NeRF as a *flexible* 3D model that can be enriched with volu-

metric descriptors to establish 2D-3D matches between an image and the scene. CrossFire [41] augments the base NeRF model with an additional 3D feature prediction branch. This feature is supervised to match 2D features extracted by an image backbone. NeRF-Loc [37] performs feature matching by utilizing a combination of features extracted from a generalizable NeRF and projected multi-view image features as 3D point features. However, both methods require training NeRF jointly with the matching task, prohibiting the usage of pre-built NeRF scenes.

In contrast, our work treats NeRF as the primary scene representation in visual localization without re-training and modifications. Specifically, we focus on a crucial component – NeRF features – and (i) demonstrate their inherent capability in effectively supporting feature matching. (ii) We introduce, NeRFMatch, a matching transformer that aligns 2D image features with 3D NeRF features and a minimal version of it to facilitate real-time applications. (iii) We present two options for performing pose refinement on top of NeRFMatch and conduct detailed analysis on their refinement effectiveness. (iv) We use our NeRFMatch and pose refinement modules to perform hierarchical NeRF localization which achieves competitive localization performance on Cambridge Landmarks [31]. Based on our experiments, we point out future work in needs to improve our indoor localization performance. Our research paves the path towards localization leveraging NeRF as the sole representation of the scene.

## 2 Related work

### 2.1 Visual Localization

**Structure-based localization.** Such methods [12, 28, 36, 48, 51, 57] first estimate 2D-3D correspondences between a query image and the 3D points in the scene and then deploy a Perspective-n-Point (PnP) solver [23, 29, 33] to compute the query camera pose. Image retrieval [1, 4, 24, 60] is usually applied in advance to coarsely localize visible scene structure to a query image [48, 57]. Visual features [14, 19, 21, 49, 56, 63, 73] are often extracted from a database of scene images to represent 3D features and matched against the query image features extracted using the same algorithm to obtain 2D-3D matches. To optimize the inference runtime, 3D descriptors are cached with the scene model at the cost of high storage demand and challenging map updates. To relieve the burden coming from the need of storing visual descriptors, GoMatch [72] performs geometric feature matching, yet currently being less accuracy than visual feature matching approaches. Recently, to avoid storing massive amount of scene images as well as being flexible to switch between different descriptors for 2D-3D matching, MeshLoc [44] propose to use 3D meshes as a dense 3D model. Compared to MeshLoc, we also pursuit a dense scene representation via NeRF [40], which not only store scene images compactly but also provides free per-3D-point NeRF descriptors for direct 2D-to-3D matching.

**End-to-end Learned Localization.** APR methods [5, 11, 30, 31, 53, 62] directly regress a camera pose from a query image, while being lightweight by

encoding scene information within a single model, they are currently less accurate than approaches based on 2D-3D matching [52]. In contrast, SCR methods [6, 8–10, 35, 54, 67] perform implicit 2D-3D matching via directly regressing 3D scene coordinates from a query image. Similar to APR, they learn to encode the scene geometry within their own network parameters [6, 8–10, 35] but they are limited by the model capacity to memorize large-scale scenes. Recently, scene agnostic SCR methods [59, 67] have been proposed to scale up to larger scenes by decoupling the scene representation from the learned matching function.

## 2.2 NeRF in Localization

iNeRF [68] directly inverts a NeRF model to refine a camera pose initialization by iteratively optimizing the photometric difference between the rendered image and the query image. However, it requires hundreds of iterations to converge and thus is not directly applicable to real-world visual localization. LENS [42] leverages NeRF as a novel view synthesizer (NVS) to augment the image database for pose regression training. Similar to LENS [42], NeRF-SCR [15] augments SCR training samples using RGB-D images rendered from a NeRF model based on an uncertainty-guided novel view selection. DirectPN [18] incorporates NeRF to provide photometric consistency supervision for pose regression where it minimizes the color difference between the query image and the image rendered at a predicted pose. DFNet [17] extends this idea to measure the consistency in the feature space showing boosted localization performance. NeFes [16] follows iNeRF to use NeRF as an offline pose refinement module on top of DFNet. It distills pre-trained DFNet feature into a NeRF model and directly renders it for pose optimization. Different from NeFes, we directly use 3D viewpoint-invariant feature learned during a standard NeRF training and validating its potential to deliver highly-accurate localization performance. In addition, our model is able to perform localization in a scene-agnostic (multi-scene) setting while DFNet and NeFes is scene-dependent.

The most relevant works to ours are NeRFLoc [37] and CrossFire [41] as they also establish explicit 2D-3D matches with features rendered from NeRF. NeRFLoc proposes a generalizable NeRF that is conditioned on a set of reference images and reference depths to output descriptors for 3D points by fusing multi-view image features, while CrossFire lifts an instant-NGP [43] model to directly outputs feature descriptors for 3D points. In contrast to their methods that both require to train their customized scene model together with the matching model, our NeRFMatch directly learn to align image feature with pre-trained NeRF features, which allows us to directly benefit from the on-going advancement in the typical NeRF research.

## 3 NeRF-based Localization

In this work we explore the capability of NeRF features, tasked with view synthesis, to offer precise 2D-3D correspondences for addressing visual localization.

To this end, we introduce our NeRF-based localization pipeline, which adheres to the general steps of standard structure-based localization approaches [48,57]. We provide an overview of the localization pipeline in Sec. 3.1, followed by an explanation on how we utilize NeRF as a scene representation for localization in Sec. 3.2. Afterwards, we detail our iterative pose refinement component in Sec. 3.3. Finally, we delve into the specific challenge of matching images to NeRF features using our newly proposed NeRFMatch model in Sec. 4.

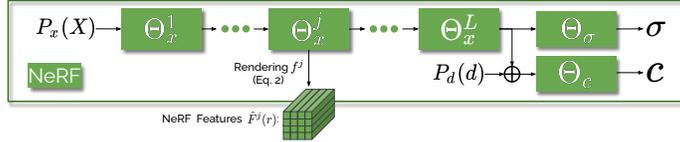
### 3.1 Localization Pipeline

Given a query image and a scene represented by a pre-trained NeRF model associated with a list of pre-cached reference poses, *e.g.*, the poses used to train the NeRF, our localization pipeline involves three steps to localize the query as shown in Fig. 1. (i) We first apply image retrieval [1,60], which extracts visual descriptors from the query image and the reference images (synthesized by NeRF at all reference poses, *cf.* Sec. 5.1) to find the nearest neighbour to the query based on descriptor distances. This step efficiently narrows down the pose search space to the vicinity of the reference pose. (ii) We use the reference pose to render a set of 3D points with associated NeRF descriptors (*cf.* Sec. 3.2) and feed it together with the query image into our NeRFMatch (*cf.* Sec. 4) to predict the 2D-3D matches between the query image and 3D points, from which we estimate the absolute camera pose of the query image via a PnP solver [23,29,33]. (iii) We further use our pose refinement module (*cf.* Sec. 3.3) to improve the pose estimation iteratively. Note that this step is optional considering the trade-off between accuracy and runtime efficiency.

### 3.2 NeRF for Localization

**NeRF architecture.** NeRF [40] models a 3D scene with a coordinate network. Specifically, it maps a 3D point and a camera viewing direction to their corresponding volume density and RGB values. As depicted in Fig. 2, a standard NeRF model consists of two non-learnable positional encodings,  $P_x$  for 3D coordinates and  $P_d$  for viewing directions, and three trainable components: a 3D point encoder  $\Theta_x$ , a volume density decoder  $\Theta_\sigma$ , and a RGB decoder  $\Theta_c$ . More concretely, the 3D encoder consists of  $L$  layers, *i.e.*,  $\Theta_x = \Theta_x^L \circ \dots \circ \Theta_x^1$ . In this work, we are particularly interested in exploring the potential of 3D features extracted within the 3D encoder. Given a 3D point  $X \in R^3$ , we define its 3D feature extracted at  $j$ -th 3D encoder layer as  $f^j = \Theta_x^j \circ \dots \circ \Theta_x^1(P_x(X))$ , where  $\Theta_x^j$  is the  $j$ -th layer in the 3D encoder. The last layer 3D feature is next input into the density decoder to predict the volume density  $\sigma = \Theta_\sigma(f^L)$ , while the color decoder takes a positional-embedded viewing direction  $d \in R^2$  in addition to the 3D feature  $f^L$  to compute the view-dependent color  $c = \Theta_c(f^L, P_d(d))$ .

**Volumetric rendering of color.** NeRF employs the continuous scene representation to render per-ray color using a discretized volumetric rendering procedure. Given a ray  $r(t) = o + td$  emitted from the camera center  $o \in R^3$ , along a viewing direction  $d$  which intersects the image plane at pixel  $x$ , this ray is



**Fig. 2: An overview of the standard NeRF architecture.** The input consists of a scene coordinate  $X$  and ray directions  $d$ . The outputs include color  $c$ , density  $\sigma$ . We obtain intermediate features, denoted as  $f^j$ , using volumetric rendering.

sampled at  $N$  3D points between a near and far planes. The RGB color  $\hat{C}(r)$  at pixel  $x$  is calculated as:

$$\hat{C}(r) = \sum_{i=1}^N w_i c_i, \quad w_i = T_i (1 - \exp(-\delta_i \sigma_i)), \quad (1)$$

where  $\delta = t_{i+1} - t_i$  is the sampling interval,  $c_i$  and  $\sigma_i$  are the predicted color and density of  $i$ -th sampled 3D point, and  $T_i = \exp(-\sum_{j=1}^{i-1} \sigma_j \delta_j)$  [40].

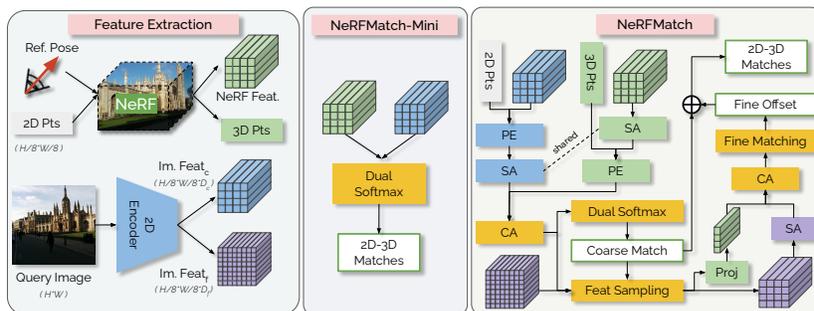
**Volumetric rendering of 3D points and features.** We investigate the capability of NeRF features to act as descriptors for 3D surface points, facilitating 2D-3D correspondence with a query image. Following the volumetric rendering process defined in Eq. (1), we define a rendered 3D (surface) point  $\hat{X}(r)$  and its associated NeRF descriptor  $\hat{F}^j(r)$  along the ray  $r$  cast through image pixel  $x$  as the weighted sum of sampled 3D points and their respective features:

$$\hat{X}(r) = \sum_{i=1}^N w_i X_i, \quad \hat{F}^j(r) = \sum_{i=1}^N w_i f_i^j. \quad (2)$$

Here,  $X_i$  is  $i$ -th sampled 3D points,  $f_i^j$  is its 3D feature extracted at  $j$ -th 3D encoder layer and  $w_i$  is the weight computed from its density prediction Eq. (1).

### 3.3 Pose Refinement

Following the matching-based pose estimation, we employ two approaches to refine the estimated camera pose. The first approach, iterative refinement, uses the estimated camera pose as a new reference for extracting NeRF features. This process involves repeating the matching procedure with the updated reference pose, incrementally enhancing the results due to the closer proximity of the reference pose, which makes the NeRF and image features more similar. While this process can be repeated multiple times, significant improvements are typically observed after the first refinement (*cf.* Sec. 5.4). Inspired by iNeRF [68], our second refinement option combines optimization and matching where we back-propagate through the frozen NeRF model to optimize an initial camera pose estimate by minimizing the photometric difference between the query image and



**Fig. 3: NeRFMatch architecture.** We present NeRFMatch as our full matching model (*rightmost*) and NeRFMatch-Mini as a light version of it (*middle*). Both models share the same feature extraction process, where we use a 2D encoder to extract image features at two resolutions and render 3D points with associated NeRF features at sampled 2D pixel locations from the reference viewpoint. The full matching uses self-attention (SA) and cross-attention (CA) with positional encodings (PE).

the current NeRF rendered image. We then use the optimized camera pose to render again 3D points and features to perform the 2D-3D matching and compute the final refined camera pose.

## 4 Image-to-NeRF Matching Network

**Overview.** To establish 2D-3D correspondences between a query image and NeRF scene points for localization, we introduce two variants of our proposed image-to-NeRF matching network, a *full* version (NeRFMatch) and a *minimal* version (NeRFMatch-Mini), as depicted in Fig. 3. The *full* version is notably more expressive, incorporating powerful attention modules [14, 49, 56] and follows a coarse-to-fine matching paradigm [14, 56, 73]. While it delivers more accurate matches it is computationally more expensive (*cf.* Sec. 5.3). Consequently, we also propose a *minimal* version of our method that focuses on learning good features for matching, removing the need to learning the matching function itself. Both models comprise a feature extraction module that encodes both the query image and 3D scene points into a feature space, and a matching module that aligns these two feature sets to determine the 2D-3D correspondences. We further detail their architectural designs and supervision in the subsequent subsections.

### 4.1 NeRFMatch-Mini

**Image encoding.** Given a query image  $I \in \mathbb{R}^{H \times W \times 3}$ , we use a CNN encoder to extract a coarse-level feature map  $F_m^c \in \mathbb{R}^{N_m \times D^c}$  with  $N_m = \frac{H}{8} \times \frac{W}{8}$  and a fine-level feature map  $F_m^f \in \mathbb{R}^{\frac{H}{2} \times \frac{W}{2} \times D^f}$ .

**NeRF feature encoding.** Given the reference pose of the query image found by image retrieval (*cf.* Sec. 3.1), we use NeRF to obtain a set of scene 3D points  $X_s \in \mathbb{R}^{N_s \times 3}$  and their associating NeRF features  $F_s \in \mathbb{R}^{N_s \times D^c}$  (*cf.* Sec. 3.2). To make the matching module memory managable, we consider NeRF features  $F_s \in \mathbb{R}^{N_s \times D^c}$  rendered along rays originated at center pixels of every  $8 \times 8$  image patch in the reference view, *i.e.*,  $N_s = \frac{H}{8} \times \frac{W}{8}$ .

**Dual-softmax matching.** After the feature extraction, NeRFMatch-Mini directly matches the NeRF feature map  $F_s$  against the coarse-level image feature map  $F_m^c$  with a non-learnable dual-softmax matching function adopted from [56]. Specifically, we compute pair-wise cosine similarities between the two feature maps followed by a dual-softmax operation to obtain the matching score matrix  $S \in \mathbb{R}^{N_m \times N_s}$ . Finally, we extract mutual matches based on the association scores which gives us the predicted matches.

## 4.2 NeRFMatch

As depicted in Fig. 3, NeRFMatch shares the same feature extraction processes as NeRFMatch-Mini, but learns an attention-based matching module which follows the coarse-to-fine paradigm of image-to-image feature matching [14, 56, 73]: (i) firstly, we identify 3D-to-image patch matches using a coarse matching module, (ii) secondly, we refine to pixel accuracy with a fine matching module. We describe the detailed matching module architecture in the following paragraphs.

**Coarse-level matching.** Equipped with coarse image features  $F_m^c$  that represent image local patches and raw NeRF features  $F_s$  that represent individual 3D scene points, we first apply 2D positional encoding [13, 56] to equip the image features with positional information. We next enrich the global contextual information within each domain by applying several self-attention blocks. We share the self-attention weights between image features and NeRF features to help bring the features from two different domains into a common embedding space for matching. After self-attention, we enhance the 3D information explicitly by concatenating each NeRF feature with its positional-encoded 3D points using NeRF positional encoding [40]. Afterwards, we feed those features into a cross-attention layer to enable cross-domain interactions. We use the same dual-softmax matching function introduced above (*cf.* Sec. 4.1) to obtain coarse matches  $\mathcal{M}_c = \{(i, j) | i \in (0, N_m - 1), j \in (0, N_s - 1)\}$  where  $i, j$  are the indices of the image and point features.

**Fine-level matching.** For each coarse match  $m_c = (i, j)$  that we extract, we start by gathering its high-resolution image patch feature  $F_m^f(i) \in \mathbb{R}^{w \times w \times D^f}$  from the fine-level feature map, centered around the corresponding location of the match. Inside each of these image patches, a self-attention block is applied to spread contextual information throughout the patch. Next, for every 3D point feature, we take the cross-attended feature obtained from the coarse matching, and process it through a linear layer to adjust its feature dimension from  $D^c$  to  $D^f$ . In the subsequent step, we align the 3D feature with its corresponding local feature map. This alignment [56] produces a heatmap, which depicts the likelihood of the 3D point  $j$  matching with each pixel in the vicinity of image pixel

*i.* To obtain the exact, fine match, we compute the expected value across this heatmap’s probability distribution. The final, refined matches obtained through this process are denoted as  $\mathcal{M}_f$ .

### 4.3 Supervision

**Ground-truth matches.** We project the rendered 3D scene points  $X_s$  onto the query image using the ground-truth query camera calibration, which gives their precise 2D projections  $x_s$  at the resolution in which we want to supervise the fine matches. To compute the ground-truth coarse association matrix  $M_{gt}$  which is a binary mask, for each 3D point  $j$ , we assign it to its belonging  $i$ -th  $8 \times 8$  local patch in the query image. We set the association value  $M_{gt}(i, j) = 1$  if a 3D point  $j$  finds its 2D patch  $i$  within the query image boundary. Notice, one image patch can be assigned to multiple 3D points yet each 3D point has at most one 2D match.

**Losses.** To supervise the coarse matching, we apply the log loss [56] to increase the dual-softmax probability at the ground-truth matching locations in  $M_{gt}$ . The coarse matching loss  $L_c$  is defined as:

$$L_c = -\frac{1}{M_{gt}} \sum_{(i,j) \in \mathcal{M}_{gt}} \log(S(i, j)). \quad (3)$$

To compute the fine matching loss, for a 3D point  $X_j$  with ground-truth fine match  $x_j$ , we supervise its predicted fine match  $\tilde{x}_j$  by minimizing its pixel distance to the ground-truth match. Following [56, 63], we compute the total variance  $\sigma^2(j)$  of the corresponding heatmap and minimize the weighted loss function:

$$L_f = \frac{1}{M_f} \sum_{(i,j) \in \mathcal{M}_f} \frac{1}{\sigma^2(i)} \|\tilde{x}_j - x_j\|_2. \quad (4)$$

The NeRFMatch-Mini model is supervised with only the coarse loss  $L_c$  while for NeRFMatch model is supervised with the sum of coarse and fine matching losses  $L_c + L_f$ .

## 5 Experiments

**Datasets.** Cambridge Landmarks [31] is a dataset of handheld smartphone images of 6 outdoor scenes with large exposure variations which is considered challenging for NeRF techniques. We follow previous work [6, 50] and evaluate on a subset of 5 scenes whose spatial extent ranges from  $875m^2$  to  $5600m^2$ . We also test our method on 7-Scenes [54], which is composed of RGB-D images captured in 7 unique indoor scenes whose size ranges from  $1m^3$  to  $18m^3$ . Its images contain large texture-less surfaces, motion blur, and object occlusions. For both datasets, we follow the original released training and testing splits. Following recent work [6, 7, 16], we use the more accurate SfM pose annotations for 7-Scenes rather than its original pose annotations.

**Evaluation metrics.** We report median pose errors, *i.e.*, translation error in centimeters and median rotation error in degrees. We further report localization recall that measures the percentage of queries localized with pose errors below specific thresholds, *i.e.*,  $(5cm, 5^\circ)$  for 7-Scenes. As Cambridge Landmarks has large variation in scene scales, we follow [10, 37] to use  $5^\circ$  rotation error and variable translation error thresholds, *i.e.*,  $38/22/15/35/45cm$  for King’s College/Old Hospital/Shop Facade/St. Mary’s Church/Great Court.

**Implementation details.** We use the first two blocks of ConvFormer [69] as the image backbone and initialize it with ImageNet-1K [47] pre-trained weights\*. We set feature dimensions for coarse and fine matching as  $D^c = 256$  and  $D^f = 128$ . For fine matching, we use local window size  $w = 5$  for image feature cropping. We resize query images to  $480 \times 480$  for all experiments. We train minimal and full NeRFMatch models using Adam [32] optimizer with canonical learning rate  $clr = 0.0008/0.0004$  and batch size  $cbs = 16$  for 30/50 epochs accordingly. We decay the learning rate based on the cosine annealing scheduling [38]. Our models are trained on 8 Nvidia V100 GPUs (16/32GB). A MipNeRF model is 5.28 MB in size. Given a camera pose, it takes 141 milliseconds to render 3600 3D points with its features on a single 16GB Nvidia V100 GPU. Our NeRFMatch-Mini and NeRFMatch models are 42.8/50.4 MB in size and require 37/157ms to run a forward pass at  $480 \times 480$  image resolution. We provide NeRF implementation details in the supplementary material.

## 5.1 Localization Evaluation

**Baselines.** We first compare our proposed NeRFMatch and NeRFMatch-Mini against common visual localization approaches on both indoor and outdoor datasets. We split the methods into the *end-to-end* category including APR [6, 17, 42, 53] and SCR [6, 10, 35] methods, and the *hierarchical* category [37, 41, 48, 50, 57–59, 67] where methods rely on an extra image retrieval step to coarsely localize the region in the scene. In addition, we specify the underlying scene representation used for localization at *test time*.

**Inference settings.** We use *top-1/10* reference poses for Mini/Full NeRFMatch models for outdoor and *top-1* for indoor and apply the best pose refinement determined in Sec. 5.4. We provide more details in supplementary.

**Results on Cambridge Landmarks.** As shown in Tab. 1, our *minimal* version despite being lightweight, is able to achieve comparative results w.r.t. most of the SCR methods and surpass all APR methods. With a more advanced attention-based hierarchical matching function, our *full* model achieves the competitive results among all methods. Our experiments fully demonstrate that the NeRF inner features learned via view synthesis are discriminative 3D representations for 2D-3D matching.

**Results on 7-Scenes.** As shown in Tab. 2, we are on-par with the best APR method (the *upper* rows), NeFeS [16], while we are less accurate than SCR and

---

\* The weights can be downloaded from [huggingface.co/timm/convformer\\_b36\\_sail\\_in1k\\_384](https://huggingface.co/timm/convformer_b36_sail_in1k_384)

**Table 1: Outdoor localization on Cambridge Landmarks [31].** We report per-scene median rotation and position errors in  $(cm, ^\circ)$  and its average across scenes.

Method	Scene Repres.	Cambridge Landmarks - Outdoor						
		Kings	Hospital	Shop	StMary	Court	Avg.Med ↓	
End-to-End	MS-Trans. [53]	APR Net.	83/1.5	181/2.4	86/3.1	162/4	-	-
	DFNet [17]	APR Net.	73/2.4	200/3	67/2.2	137/4	-	-
	LENS [42]	APR Net.	33/0.5	44/0.9	27/1.6	53/1.6	-	-
	NeFeS [16]	APR+NeRF	37/0.6	55/0.9	14/0.5	32/1	-	-
	DSAC* [10]	SCR Net.	15/0.3	21/0.4	5/0.3	13/0.4	49/0.3	20.6/0.3
	HACNet [35]	SCR Net.	18/0.3	19/0.3	6/0.3	9/0.3	28/0.2	16/0.3
Hierarchical	ACE [6]	SCR Net.	28/0.4	31/0.6	5/0.3	18/0.6	43/0.2	25/0.4
	SANet [67]	3D+RGB	32/0.5	32/0.5	10/0.5	16/0.6	328/2.0	83.6/0.8
	DSM [58]	SCR Net.	19/0.4	24/0.4	7/0.4	12/0.4	44/0.2	21.2/0.4
	NeuMap [59]	SCode+RGB	14/0.2	19/0.4	6/0.3	17/0.5	6/0.1	12.4/0.3
	InLoc [57]	3D+RGB	46/0.8	48/1.0	11/0.5	18/0.6	120/0.6	48.6/0.7
	HLoc [48]	3D+RGB	12/0.2	15/0.3	4/0.2	7/0.2	16/0.1	10.8/0.2
	PixLoc [50]	3D+RGB	14/0.2	16/0.3	5/0.2	10/0.3	30/0.1	15/0.2
	CrossFire [41]	NeRF+RGB	47/0.7	43/0.7	20/1.2	39/1.4	-	-
	NeRFLoc [37]	NeRF+RGBD	11/0.2	18/0.4	4/0.2	7/0.2	25/0.1	13/0.2
	NeRFMatch-Mini	NeRF+RGB	19.0/0.3	30.2/0.6	10.3/0.5	11.3/0.4	29.1/0.2	20.0/0.4
	NeRFMatch	NeRF+RGB	13.0/0.2	19.4/0.4	8.5/0.4	7.9/0.3	17.5/0.1	13.3/0.3
	NeRFMatch	NeRF	12.7/0.2	20.7/0.4	8.7/0.4	11.3/0.4	19.5/0.1	14.6/0.3

**Table 2: Indoor localization on 7-Scenes [54].** We report per-scene median rotation and position errors in  $(cm, ^\circ)$  and their average across scenes, along with averaged localization recall.

Method	Scene Repres.	7-Scenes - SfM Poses - Indoor								
		Chess	Fire	Heads	Office	Pump.	Kitchen	Stairs	Avg.Med ↓	Avg.Recall ↑
MS-Trans. [53]	APR Net.	11/6.4	23/11.5	13/13	18/8.1	17/8.4	16/8.9	29/10.3	18.1/9.5	-
DFNet [17]	APR Net.	3/1.1	6/2.3	4/2.3	6/1.5	7/1.9	7/1.7	12/2.6	6.4/1.9	-
NeFeS [16]	APR+NeRF	2/0.8	2/0.8	2/1.4	2/0.6	2/0.6	2/0.6	5/1.3	2.4/0.9	-
DSAC* [10]	SCR Net.	0.5/0.2	0.8/0.3	0.5/0.3	1.2/0.3	1.2/0.3	0.7/0.2	2.7/0.8	1.1/0.3	97.8
ACE [6]	SCR Net.	0.7/0.5	0.6/0.9	0.5/0.5	1.2/0.5	1.1/0.2	0.9/0.5	2.8/1.0	1.1/0.6	97.1
DVLAD+R2D2 [45, 60]	3D+RGB	0.4/0.1	0.5/0.2	0.4/0.2	0.7/0.2	0.6/0.1	0.4/0.1	2.4/0.7	0.8/0.2	95.7
HLoc [48]	3D+RGB	0.8/0.1	0.9/0.2	0.6/0.3	1.2/0.2	1.4/0.2	1.1/0.1	2.9/0.8	1.3/0.3	95.7
NeRFMatch-Mini	NeRF+RGB	1.6/0.5	1.5/0.6	1.4/0.9	3.6/1.0	3.5/0.9	1.7/0.5	8.5/2.1	3.1/0.9	74.4
NeRFMatch	NeRF+RGB	0.9/0.3	1.1/0.4	1.4/1.0	3.0/0.8	2.2/0.6	1.0/0.3	9.0/1.5	2.7/0.7	78.2
NeRFMatch	NeRF	0.9/0.3	1.1/0.4	1.5/1.0	3.0/0.8	2.2/0.6	1.0/0.3	10.1/1.7	2.8/0.7	78.4

visual matching methods whose accuracy on 7-Scenes is close to saturation (the *middle* rows). Our hypothesis is that those methods are able to benefit from the dense distribution of frames in the sequence. However, we have to limit our NeRF training to 900 training frames (loaded at once into memory) per scene for efficient training despite thousands of frames are available.

**NeRF-only localization.** We further push our method to using NeRF as the *only* scene representation for localization, which means we would no longer need access to a real image database as part of the scene representation, which requires a significantly bigger storage than a single NeRF model. For this purpose, we propose to perform image retrieval on synthesized images rendered by our NeRF model. This entails a small decrease in performance on Cambridge in translation error. This slight degradation in performance can be attributed to the increased complexity of these scenes for NeRF. This assertion is supported by the observation that, when switching to image retrieval on synthesized images, our model demonstrates almost no change in performance on the indoor

**Table 3: NeRF feature ablation on Cambridge [31].** We train NeRFMatch-Mini with different 3D features and compare their localization performance.

Metrics	Pt3D	Pe3D	$f^1$	$f^2$	$f^3$	$f^4$	$f^5$	$f^6$	$f^7$
Med. Translation (cm, ↓)	458.0	34.3	28.7	28.4	<b>27.9</b>	28.3	28.3	30.2	61.3
Med. Rotation (°, ↓)	6.5	0.6	<b>0.5</b>	<b>0.5</b>	<b>0.5</b>	<b>0.5</b>	<b>0.5</b>	<b>0.5</b>	1.3
Localize Recall. (% , ↑)	0.7	51.4	58.6	<b>59.4</b>	<b>59.2</b>	56.9	57.7	53.0	38.8

7-Scenes dataset. Our results open the door to a localization pipeline where only a NeRF model is needed.

In the following experiments, we conduct ablation studies to thoroughly understand the different components of our method. We conduct all ablations on Cambridge Landmarks.

## 5.2 NeRF Feature Ablation

**Experimental setup.** As our first ablation, we investigate the potential of NeRF features in 2D-3D matching. We consider several different types of 3D features including the raw 3D point coordinates (Pt3D), positional encoded 3D points (Pe3D), and NeRF inner features output from all intermediate layers  $f^j$  with  $j \in [1, 7]$ , as shown in Fig. 2. NeRF features have the same dimension as the image backbone features, and thus are directly ready for 2D-3D matching. For 3D point coordinates and positional encoded 3D points as 3D features, we use a simple linear layer to lift them to the image feature dimension, and we train them together with the image backbone on feature matching. To fully focus on the influence of different 3D features, we train NeRFMatch-Mini models with different 3D features and conduct matching without pose refinement. We evaluate the quality of 2D-3D matches for localization and report the median pose errors and localization recall in Tab. 3.

**Results.** We show that directly matching image features with lifted 3D coordinate features does not yield accurate results. This significantly improves when using a NeRF positional encoding layer [40] on top of the raw 3D coordinates. Notably, the Pe3D feature corresponds to the input for NeRF. We further show that all first six layer NeRF features are better than Pe3D in aligning image features to 3D features, showing that matching benefits from the view synthesis training. The only exception is that the last layer feature  $f^7$ , which produces features that are less discriminative for matching and are more focused on the NeRF goal of predicting density and RGB values. Among layer features, we found the middle layer  $f^3$  to be slightly better than the others and choose it to be the default NeRF feature we use for our NeRFMatch models.

## 5.3 NeRFMatch Ablation

**Architecture ablation.** Next, we study the influence of different image backbones and matching functions on our matching model in Tab. 4. We first compare two convolution backbones for image feature encoding, *i.e.*, ResNet34 [25] and

**Table 4: NeRFMatch architecture ablation on Cambridge Landmarks [31].** We report averaged median pose error in  $(cm, ^\circ)$  and localization recall.

Model Name	Backbone	Pretrain Backbone	Matcher	Training Scenes	Avg.Med $(cm/^\circ) \downarrow$	Avg.Recall $(\%) \uparrow$	Model Size $(MB) \downarrow$	Runtime $(ms) \downarrow$
-	ResNet34	✓	Minimal	Per-Scene	32.7/0.6	52.0	<b>32.8</b>	<b>23</b>
-	ConvFormer	✗	Minimal	Per-Scene	34.8/0.6	50.7	42.8	37
NeRFMatch-Mini	ConvFormer	✓	Minimal	Per-Scene	27.9/0.5	59.2	42.8	37
NeRFMatch-Mini (MS)	ConvFormer	✓	Minimal	Multi-Scene	30.8/0.5	53.6	42.8	37
NeRFMatch	ConvFormer	✓	Full	Per-Scene	<b>16.5/0.3</b>	<b>71.3</b>	50.4	157
NeRFMatch (MS)	ConvFormer	✓	Full	Multi-Scene	22.0/0.4	65.2	50.4	157

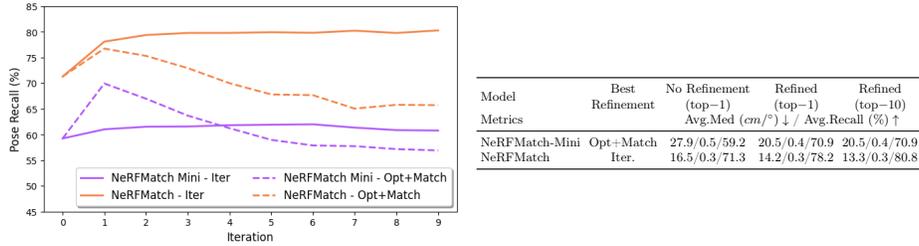
ConvFormer [69]. Our method demonstrates improved performance with the latter, which is consistent with their performance on image classification [47]. We also observe that large-scale ImageNet [47] pre-training provides a better starting point for extracting suitable image features for matching, leading to increased localization accuracy compared to training from scratch. We further show that for the same backbone, a more advanced attention matching function with a coarse-to-fine design is crucial for accurate matching and significantly improves localization accuracy, albeit at the cost of increased runtime.

**Training ablation.** In addition to architecture choices, we also examine the influence of different training settings, *i.e.*, training per-scene (default) and training multi-scenes (all scenes within a dataset). Despite NeRF features being trained per-scene, we surprisingly find that both minimal and full NeRFMatch models can be trained to handle multi-scene (MS) localization with only a slight decrease in accuracy w.r.t. the per-scene model. This is a similar finding as recent scene-agnostic SCR methods [58, 59, 67], that extend per-scene SCR to multi-scenes by conditioning SCR on scene-specific 3D points. While scene-agnostic SCR learns to regress directly the 3D features in the form of xyz coordinates, our models learn to find a common ground between image features and NeRF features for matching.

#### 5.4 Pose Refinement

After conducting matching architecture ablation, we investigate different refinement methods to further increase pose accuracy. We examine two approaches: (i) an iterative approach, where we re-run matching with the last computed estimation as reference pose, and (ii) an optimization approach, where we optimize a reference pose through frozen NeRF weights using a photometric loss and then run the final matching. As shown in Fig. 4, both methods show an improvement over the initial estimate. We assess both the NeRFMatch model and its minimal version (Mini) trained per-scene.

**Iterative refinement for computational efficiency.** In the NeRFMatch setting, both refinement approaches show similar early results, since the initial pose estimate is relatively close to the solution, but the iterative approach is more stable over time. Additionally, it is worth noting that the optimization approach incurs a higher computational cost compared to the iterative approach. The runtime for the optimization refinement excluding the matching step (shown



**Fig. 4: Refinement ablation on Cambridge Landmarks [31].** On the left side, we depict the average recall for optimization-based (Opt+Match) and iterative (Iter) refinement approaches across multiple iterations. We provide results for both the NeRF-Match and its minimal setting. On the right side, we report averaged median pose error in ( $cm/^\circ$ ) and localization recall with the best refinement configurations.

in Tab. 4) is  $398.4ms$  for a single optimization step and subsequent rendering. In contrast, for a single step of iterative refinement, the runtime amounts to  $141.2ms$ . Given the better performance and the quicker computational time, we opt for the iterative approach as the default refinement for NeRFMatch.

**Optimization refinement for large pose corrections.** In comparison to the iterative method, the optimization approach significantly benefits the minimal matching model. This is because the initial query pose estimation is notably distant compared to the one given by NeRFMatch, and subsequent iterations only result in incremental refinements. The optimization approach takes a more substantial step towards the query pose, and achieving a more optimal reference pose results in improved query pose estimation. We have already seen that the optimization refinement incurs higher computational time, but it also depends on a learning rate schedule, which is sensitive to configuration. Through empirical analysis, we select  $1 \times 10^{-3}$  as the initial learning rate and apply a cosine annealing learning rate schedule. We set the optimization approach as the default refinement for NeRFMatch-Mini.

## 6 Conclusion and Limitations

In this work, we have taken initial steps towards leveraging NeRF as the primary representation for the task of camera localization. To achieve this, we have thoroughly examined the performance of NeRF features in the localization task, considering various architectural designs, feature extraction from different encoder layers, and diverse training configurations. Additionally, we have demonstrated that NeRF can remove the need for the original image set for coarse localization. Our results suggest that NeRF features are highly effective for 2D-3D matching.

While NeRFMatch marks a significant step towards comprehensive localization using NeRF, it also highlights several limitations that necessitate further research. Specifically, we observe a noticeable performance gap when applying our method to indoor 7-Scenes dataset.

## References

1. Arandjelovic, R., Gronat, P., Torii, A., Pajdla, T., Sivic, J.: Netvlad: Cnn architecture for weakly supervised place recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5297–5307 (2016)
2. Arth, C., Wagner, D., Klopschitz, M., Irschara, A., Schmalstieg, D.: Wide area localization on mobile phones. In: 2009 8th IEEE International Symposium on Mixed and Augmented Reality. pp. 73–82. IEEE (2009)
3. Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5855–5864 (2021)
4. Berton, G., Masone, C., Caputo, B.: Rethinking visual geo-localization for large-scale applications. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4878–4888 (2022)
5. Blanton, H., Greenwell, C., Workman, S., Jacobs, N.: Extending absolute pose regression to multiple scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 38–39 (2020)
6. Brachmann, E., Cavallari, T., Prisacariu, V.A.: Accelerated coordinate encoding: Learning to relocalize in minutes using rgb and poses. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5044–5053 (2023)
7. Brachmann, E., Humenberger, M., Rother, C., Sattler, T.: On the limits of pseudo ground truth in visual camera re-localisation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6218–6228 (2021)
8. Brachmann, E., Krull, A., Nowozin, S., Shotton, J., Michel, F., Gumhold, S., Rother, C.: Dsac-differentiable ransac for camera localization. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6684–6692 (2017)
9. Brachmann, E., Rother, C.: Learning less is more-6d camera localization via 3d surface regression. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4654–4662 (2018)
10. Brachmann, E., Rother, C.: Visual camera re-localization from rgb and rgb-d images using dsac. *IEEE transactions on pattern analysis and machine intelligence* **44**(9), 5847–5865 (2021)
11. Brahmabhatt, S., Gu, J., Kim, K., Hays, J., Kautz, J.: Geometry-aware learning of maps for camera localization. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2616–2625 (2018)
12. Camposeco, F., Cohen, A., Pollefeys, M., Sattler, T.: Hybrid scene compression for visual localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7653–7662 (2019)
13. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: European conference on computer vision. pp. 213–229. Springer (2020)
14. Chen, H., Luo, Z., Zhou, L., Tian, Y., Zhen, M., Fang, T., Mckinnon, D., Tsin, Y., Quan, L.: Aspanformer: Detector-free image matching with adaptive span transformer. In: European Conference on Computer Vision. pp. 20–36. Springer (2022)
15. Chen, L., Chen, W., Wang, R., Pollefeys, M.: Leveraging neural radiance fields for uncertainty-aware visual localization. *arXiv preprint arXiv:2310.06984* (2023)

16. Chen, S., Bhalgat, Y., Li, X., Bian, J., Li, K., Wang, Z., Prisacariu, V.A.: Refinement for absolute pose regression with neural feature synthesis. arXiv preprint arXiv:2303.10087 (2023)
17. Chen, S., Li, X., Wang, Z., Prisacariu, V.A.: Dfnet: Enhance absolute pose regression with direct feature matching. In: European Conference on Computer Vision. pp. 1–17. Springer Nature Switzerland Cham (2022)
18. Chen, S., Wang, Z., Prisacariu, V.: Direct-posenet: Absolute pose regression with photometric consistency. In: 2021 International Conference on 3D Vision (3DV). pp. 1175–1185. IEEE (2021)
19. DeTone, D., Malisiewicz, T., Rabinovich, A.: Superpoint: Self-supervised interest point detection and description. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 224–236 (2018)
20. Dong, Z., Zhang, G., Jia, J., Bao, H.: Keyframe-based real-time camera tracking. In: 2009 IEEE 12th international conference on computer vision. pp. 1538–1545. IEEE (2009)
21. Dusmanu, M., Rocco, I., Pajdla, T., Pollefeys, M., Sivic, J., Torii, A., Sattler, T.: 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. pp. 8092–8101 (2019)
22. Fu, X., Zhang, S., Chen, T., Lu, Y., Zhu, L., Zhou, X., Geiger, A., Liao, Y.: Panoptic nerf: 3d-to-2d label transfer for panoptic urban scene segmentation. In: 2022 International Conference on 3D Vision (3DV). pp. 1–11. IEEE (2022)
23. Gao, X.S., Hou, X.R., Tang, J., Cheng, H.F.: Complete solution classification for the perspective-three-point problem. IEEE transactions on pattern analysis and machine intelligence **25**(8), 930–943 (2003)
24. Hausler, S., Garg, S., Xu, M., Milford, M., Fischer, T.: Patch-netvlad: Multi-scale fusion of locally-global descriptors for place recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14141–14152 (2021)
25. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
26. Heng, L., Choi, B., Cui, Z., Geppert, M., Hu, S., Kuan, B., Liu, P., Nguyen, R., Yeo, Y.C., Geiger, A., et al.: Project autovision: Localization and 3d scene perception for an autonomous vehicle with a multi-camera system. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 4695–4702. IEEE (2019)
27. Hu, B., Huang, J., Liu, Y., Tai, Y.W., Tang, C.K.: Nerf-rpn: A general framework for object detection in nerfs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 23528–23538 (2023)
28. Irschara, A., Zach, C., Frahm, J.M., Bischof, H.: From structure-from-motion point clouds to fast location recognition. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. pp. 2599–2606. IEEE (2009)
29. Ke, T., Roumeliotis, S.I.: An efficient algebraic solution to the perspective-three-point problem. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7225–7233 (2017)
30. Kendall, A., Cipolla, R.: Geometric loss functions for camera pose regression with deep learning. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5974–5983 (2017)
31. Kendall, A., Grimes, M., Cipolla, R.: Posenet: A convolutional network for real-time 6-dof camera relocalization. In: Proceedings of the IEEE international conference on computer vision. pp. 2938–2946 (2015)

32. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. 3rd International Conference on Learning Representations, (ICLR) 2015 (2015)
33. Kneip, L., Scaramuzza, D., Siegwart, R.: A novel parametrization of the perspective-three-point problem for a direct computation of absolute camera position and orientation. In: CVPR 2011. pp. 2969–2976. IEEE (2011)
34. Kundu, A., Genova, K., Yin, X., Fathi, A., Pantofaru, C., Guibas, L.J., Tagliasacchi, A., Dellaert, F., Funkhouser, T.: Panoptic neural fields: A semantic object-aware neural scene representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12871–12881 (2022)
35. Li, X., Wang, S., Zhao, Y., Verbeek, J., Kannala, J.: Hierarchical scene coordinate classification and regression for visual localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11983–11992 (2020)
36. Li, Y., Snavely, N., Huttenlocher, D.P.: Location recognition using prioritized feature matching. In: European conference on computer vision. pp. 791–804. Springer (2010)
37. Liu, J., Nie, Q., Liu, Y., Wang, C.: Nerf-loc: Visual localization with conditional neural radiance field. 2023 IEEE International Conference on Robotics and Automation (ICRA) (2023)
38. Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983 (2016)
39. Maggio, D., Abate, M., Shi, J., Mario, C., Carlone, L.: Loc-nerf: Monte carlo localization using neural radiance fields. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 4018–4025. IEEE (2023)
40. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: European Conference on Computer Vision. pp. 405–421. Springer (2020)
41. Moreau, A., Piasco, N., Bennehar, M., Tsishkou, D., Stanciulescu, B., de La Fortelle, A.: Crossfire: Camera relocalization on self-supervised features from an implicit representation. arXiv preprint arXiv:2303.04869 (2023)
42. Moreau, A., Piasco, N., Tsishkou, D., Stanciulescu, B., de La Fortelle, A.: Lens: Localization enhanced by nerf synthesis. In: Conference on Robot Learning. pp. 1347–1356. PMLR (2022)
43. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. ACM Transactions on Graphics (ToG) **41**(4), 1–15 (2022)
44. Panek, V., Kukulova, Z., Sattler, T.: Meshloc: Mesh-based visual localization. In: European Conference on Computer Vision. pp. 589–609. Springer (2022)
45. Revaud, J., De Souza, C., Humenberger, M., Weinzaepfel, P.: R2d2: Reliable and repeatable detector and descriptor. Advances in neural information processing systems **32** (2019)
46. Rosinol, A., Leonard, J.J., Carlone, L.: Nerf-slam: Real-time dense monocular slam with neural radiance fields. arXiv preprint arXiv:2210.13641 (2022)
47. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. International journal of computer vision **115**, 211–252 (2015)
48. Sarlin, P.E., Cadena, C., Siegwart, R., Dymczyk, M.: From coarse to fine: Robust hierarchical localization at large scale. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12716–12725 (2019)

49. Sarlin, P.E., DeTone, D., Malisiewicz, T., Rabinovich, A.: Superglue: Learning feature matching with graph neural networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4938–4947 (2020)
50. Sarlin, P.E., Unagar, A., Larsson, M., Germain, H., Toft, C., Larsson, V., Pollefeys, M., Lepetit, V., Hammarstrand, L., Kahl, F., et al.: 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. pp. 3247–3257 (2021)
51. Sattler, T., Leibe, B., Kobbelt, L.: Efficient & effective prioritized matching for large-scale image-based localization. *IEEE transactions on pattern analysis and machine intelligence* **39**(9), 1744–1756 (2016)
52. Sattler, T., Zhou, Q., Pollefeys, M., Leal-Taixe, L.: Understanding the limitations of cnn-based absolute camera pose regression. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3302–3312 (2019)
53. Shavit, Y., Ferens, R., Keller, Y.: Learning multi-scene absolute pose regression with transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2733–2742 (2021)
54. Shotton, J., Glocker, B., Zach, C., Izadi, S., Criminisi, A., Fitzgibbon, A.: Scene coordinate regression forests for camera relocalization in rgb-d images. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2930–2937 (2013)
55. Sucar, E., Liu, S., Ortiz, J., Davison, A.J.: imap: Implicit mapping and positioning in real-time. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6229–6238 (2021)
56. Sun, J., Shen, Z., Wang, Y., Bao, H., Zhou, X.: Loftr: Detector-free local feature matching with transformers. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8922–8931 (2021)
57. Taira, H., Okutomi, M., Sattler, T., Cimpoi, M., Pollefeys, M., Sivic, J., Pajdla, T., Torii, A.: Inloc: Indoor visual localization with dense matching and view synthesis. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7199–7209 (2018)
58. Tang, S., Tang, C., Huang, R., Zhu, S., Tan, P.: Learning camera localization via dense scene matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1831–1841 (2021)
59. Tang, S., Tang, S., Tagliasacchi, A., Tan, P., Furukawa, Y.: Neumap: Neural coordinate mapping by auto-transdecoder for camera localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 929–939 (2023)
60. Torii, A., Arandjelovic, R., Sivic, J., Okutomi, M., Pajdla, T.: 24/7 place recognition by view synthesis. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1808–1817 (2015)
61. Ventura, J., Arth, C., Reitmayr, G., Schmalstieg, D.: Global localization from monocular slam on a mobile phone. *IEEE transactions on visualization and computer graphics* **20**(4), 531–539 (2014)
62. Walch, F., Hazirbas, C., Leal-Taixe, L., Sattler, T., Hilsenbeck, S., Cremers, D.: Image-based localization using lstms for structured feature correlation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 627–637 (2017)
63. Wang, Q., Zhou, X., Hariharan, B., Snavely, N.: Learning feature descriptors using camera pose supervision. In: European Conference on Computer Vision. pp. 757–774. Springer (2020)

64. Wendel, A., Irschara, A., Bischof, H.: Natural landmark-based monocular localization for mavs. In: 2011 IEEE International Conference on Robotics and Automation. pp. 5792–5799. IEEE (2011)
65. Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V., Sridhar, S.: Neural fields in visual computing and beyond. *Computer Graphics Forum* (2022). <https://doi.org/10.1111/cgf.14505>
66. Xu, C., Wu, B., Hou, J., Tsai, S., Li, R., Wang, J., Zhan, W., He, Z., Vajda, P., Keutzer, K., et al.: Nerf-det: Learning geometry-aware volumetric representation for multi-view 3d object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 23320–23330 (2023)
67. Yang, L., Bai, Z., Tang, C., Li, H., Furukawa, Y., Tan, P.: Sanet: Scene agnostic network for camera localization. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 42–51 (2019)
68. Yen-Chen, L., Florence, P., Barron, J.T., Rodriguez, A., Isola, P., Lin, T.Y.: inerf: Inverting neural radiance fields for pose estimation. In: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 1323–1330. IEEE (2021)
69. Yu, W., Si, C., Zhou, P., Luo, M., Zhou, Y., Feng, J., Yan, S., Wang, X.: Metaformer baselines for vision. *arXiv preprint arXiv:2210.13452* (2022)
70. Zhang, Y., Tosi, F., Mattoccia, S., Poggi, M.: Go-slam: Global optimization for consistent 3d instant reconstruction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3727–3737 (2023)
71. Zhi, S., Laidlow, T., Leutenegger, S., Davison, A.J.: In-place scene labelling and understanding with implicit scene representation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15838–15847 (2021)
72. Zhou, Q., Agostinho, S., Ošep, A., Leal-Taixé, L.: Is geometry enough for matching in visual localization? In: European Conference on Computer Vision. pp. 407–425. Springer (2022)
73. Zhou, Q., Sattler, T., Leal-Taixé, L.: Patch2pix: Epipolar-guided pixel-level correspondences. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4669–4678 (2021)
74. Zhu, Z., Peng, S., Larsson, V., Xu, W., Bao, H., Cui, Z., Oswald, M.R., Pollefeys, M.: Nice-slam: Neural implicit scalable encoding for slam. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12786–12796 (2022)