

EP2P-Loc: End-to-End 3D Point to 2D Pixel Localization for Large-Scale Visual Localization

Minjung Kim Junseo Koo Gunhee Kim
Seoul National University

{minjung.kim, junseo.koo}@vision.snu.ac.kr, gunhee@snu.ac.kr

<https://github.com/minnjung/EP2P-Loc>

Abstract

Visual localization is the task of estimating a 6-DoF camera pose of a query image within a provided 3D reference map. Thanks to recent advances in various 3D sensors, 3D point clouds are becoming a more accurate and affordable option for building the reference map, but research to match the points of 3D point clouds with pixels in 2D images for visual localization remains challenging. Existing approaches that jointly learn 2D-3D feature matching suffer from low inliers due to representational differences between the two modalities, and the methods that bypass this problem into classification have an issue of poor refinement. In this work, we propose EP2P-Loc, a novel large-scale visual localization method that mitigates such appearance discrepancy and enables end-to-end training for pose estimation. To increase the number of inliers, we propose a simple algorithm to remove invisible 3D points in the image, and find all 2D-3D correspondences without keypoint detection. To reduce memory usage and search complexity, we take a coarse-to-fine approach where we extract patch-level features from 2D images, then perform 2D patch classification on each 3D point, and obtain the exact corresponding 2D pixel coordinates through positional encoding. Finally, for the first time in this task, we employ a differentiable PnP for end-to-end training. In the experiments on newly curated large-scale indoor and outdoor benchmarks based on 2D-3D-S and KITTI, we show that our method achieves the state-of-the-art performance compared to existing visual localization and image-to-point cloud registration methods.

1. Introduction

Visual localization [2, 5, 15, 22, 23, 42, 51] aims at estimating the precise 6-degree-of-freedom (DoF) camera pose (*i.e.* position and orientation) of a given 2D image within a 3D reference map. Since it can position the pose of a photo-taker in an environment without localization infor-

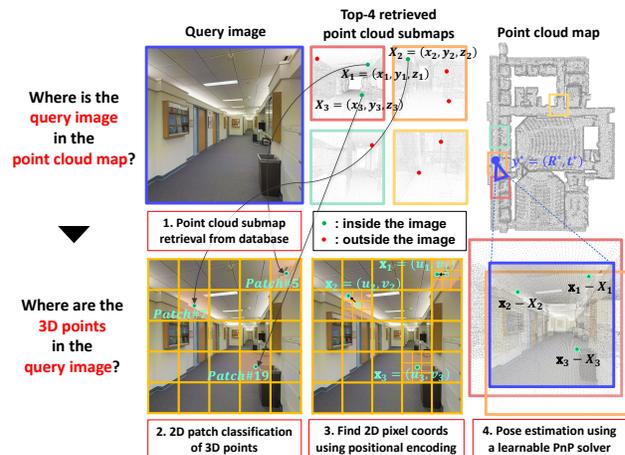


Figure 1: Visual localization of EP2P-Loc. For a query image, we (1) retrieve top-4 point cloud submaps from the database, (2) perform 2D patch classification for each 3D point in the retrieved submaps, (3) find precise 2D pixel coordinates using a positional encoding, and (4) estimate the camera pose of the image using a learnable PnP solver [11].

mation such as GPS, it has a variety of applications including autonomous driving, robot navigation, and augmented reality. Previous works use the Perspective-n-Point (PnP) algorithm [17] to compute the camera pose of a query image based on the coordinates in the 3D reference map. The 3D reference map is typically built from a collection of images using Structure-from-Motion (SfM) reconstruction [1, 24, 36, 44, 47], and the associated keypoints and descriptors are stored with the map in the form of point clouds to establish 2D-3D correspondences. Therefore, the accuracy of visual localization is highly dependent on the quality of the 3D reference map, and there are multiple challenges, including varying lighting, weather and seasonal conditions in outdoor environments, and repetitive patterns and insufficient textures in indoor environments.

With the recent advance in 3D sensors, such as RGB-D scanners and LiDARs, it is possible to generate large-scale 3D point cloud maps with no help from SfM. This could be a more practical approach compared to SfM since the point clouds are denser and more accurate. However, visual localization has difficulty in matching descriptors from 2D images with those from 3D reference point clouds due to the big discrepancy between 2D and 3D representations. This is because 2D images are represented in the form of a grid of colors and texture, while 3D point clouds are relatively sparse and lack texture information due to empty spaces in the 3D space.

To tackle this problem, several recent studies have tried to match 2D pixels of an image with 3D points of a point cloud. 2D3D-MatchNet [16] is a pioneer work that proposes a deep neural network to embed the descriptors of 2D and 3D keypoints into a shared feature space. However, it still suffers from a low inlier problem, which means that the number of predicted 2D-3D correspondences used for pose estimation is not enough, since it is difficult to mitigate the severe appearance differences between the two modalities. P2-Net [53] detects keypoints and per-point descriptors based on the RGB-D scan datasets [20, 49], but every 2D pixel must have a corresponding 3D point. Moreover, obtaining features from all 2D pixels and 3D points can be inefficient in terms of memory usage and the amount of computation for the visual localization in large-scale environments. DeepI2P [29] replaces descriptor matching with classification to reduce computation while increasing accuracy. However, it is limited for large-scale visual localization as will be shown in our experiments later, since it is challenging to find exact 2D-3D correspondences from the grid classification of each 3D point.

In this work, we propose *EP2P-Loc*, a novel approach to large-scale visual localization that prevents the low inlier problem by processing all 2D pixels in the query image and 3D points in the 3D reference maps, while reducing memory usage and search space. Figure 1 outlines the pipeline of our method at inference. Initially, the 3D point cloud reference map is divided into a set of submaps for effective candidate selection. Like a coarse-to-fine approach, we retrieve relevant 3D point cloud submaps from the database for a query image. We then perform 2D patch classification to determine to which patch of the image the 3D points in the retrieved submaps belong. Next, we find the exact 2D pixel coordinates of the 3D points using positional encoding, instead of storing features for all 2D pixels. Finally, the obtained 2D-3D correspondences pass into the PnP layer to estimate the camera pose of the image.

The key novelty of our approach lies in its ability to effectively learn features of both 2D pixels and 3D points while not only handling invisible 3D points in 2D images using our *Invisible 3D Point Removal (IPR)* algorithm but

also finding all 2D-3D correspondences without the need for keypoint detection in a coarse-to-fine manner, resulting in a higher number of inliers. Moreover, for the first time in this task, we adopt an end-to-end learnable PnP solver [11] for high-quality 6-DoF pose estimation. We demonstrate that this approach is more accurate and efficient than selecting top k pairs and manually putting them into the PnP solver. For empirical validation in large-scale indoor and outdoor environments, we establish benchmarks based on the open-source Stanford 2D-3D-Semantic (2D-3D-S) [3] and KITTI [19] datasets, and show the state-of-the-art performance compared to existing image-based visual localization methods [2, 7, 14, 15, 22, 23, 42, 43] and image-to-point cloud methods [16, 29]. Our method uses only the coordinate values of the 3D point cloud and assumes that every 3D point is not associated with a pixel in the 2D image and vice versa. Through experiments, we also demonstrate that our method is applicable to various 3D global point cloud maps generated in different ways, including 3D LiDAR data, RGB-D scans, and SfMs.

2. Related work

Visual localization. Visual localization is the problem of estimating the 6-DoF camera pose of a query image within a provided 3D reference map. Image retrieval-based approaches [2, 22, 35, 42] estimate the location by retrieving the most similar image from a geo-tagged database. The retrieved location can limit the scope of the search area within the large 3D reference map generated by SfM [1, 24, 36, 44, 47], and thus these methods allow to speed up the search in a large environment. Structure-based methods [1, 22, 24, 36, 42, 44, 45, 50, 51] extract 2D local descriptors [14, 15, 31, 41] from database images, construct the 3D reference map through SfM, and store one or more local descriptors at each 3D point. Given the query image, they extract 2D local descriptors, match local descriptors stored in the 3D reference map to find 2D-3D correspondences, and predict the 6-DoF camera pose from top k pairs using the PnP solver [27] with RANSAC [17]. Some structure-based methods [42, 51] adopt image retrieval-based methods to reduce search space for speed-up. However, these methods heavily rely on image-based features, and thus are not robust to various lighting, weather and seasonal conditions in outdoor environments, and repetitive patterns and insufficient textures in indoor environments. Some methods [5, 8, 25, 33, 46] directly regress the camera pose from the single query image, but are not competitive in terms of accuracy. PCLoc [23] and FeatLoc [5] reduce viewpoint differences between the query and the database, and estimate the pose by synthesizing new views with RGB-D images. Although they can reduce translational errors, their synthesizing modules require dense point clouds or RGB-D images, which limit the training setup or dataset selection.

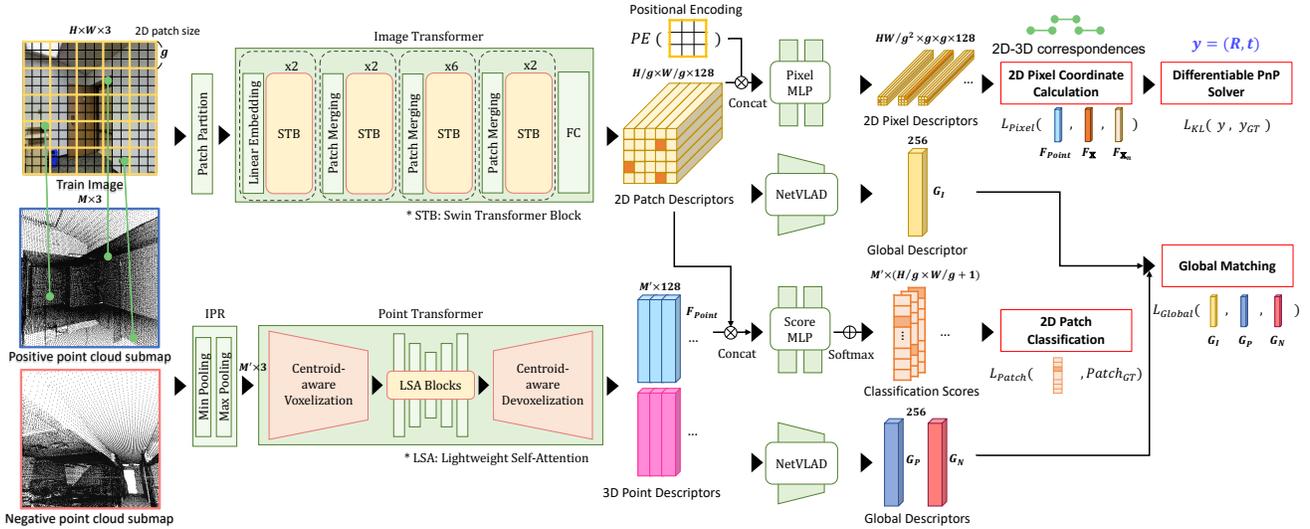


Figure 2: Overall architecture of our EP2P-Loc network.

Image-to-point cloud registration. To estimate the relative pose between an image and a point cloud, 2D3D-MatchNet [16] and LCD [38] propose deep networks to jointly learn the descriptors from the 2D image patches and 3D point cloud patches. However, they are limited in that keypoints must be set in advance, and in particular, 2D3D-MatchNet, which is the only method that aims at the same task as ours, suffers from a low inlier problem due to the difference in representation between the two modalities. 3DT-Net [54] learns 3D local descriptors that take 2D and 3D local patches as inputs while using the 2D features as auxiliary information only. Cattaneo *et al.* [10] creates a shared embedding space between 2D images and 3D point clouds adopting a teacher/student model. However, they still require 2D and 3D pre-defined keypoints, leading to inaccurate matches without consistent keypoint selection. P2-Net [53] jointly learns local descriptors and detectors for pixel and point matching, but this is applicable only when every 2D pixel has a corresponding 3D point based on RGB-D scans. DeepI2P [29] transforms the 2D-3D feature matching problem into a classification problem of whether a projected point lies within an image frustum. However, this method is limited for large-scale visual localization tasks because it cannot find exact 2D-3D correspondences and uses both an image and a point cloud as inputs to estimate the relative pose between them.

3. EP2P-Loc

Our goal is to predict the 6-DoF camera pose of a query 2D image within a 3D global reference map of a point cloud. We assume that the global map is partitioned into a set of submaps, which will be described in Section 4.1. Dur-

ing training, we are given the image $I \in \mathbb{R}^{H \times W \times 3}$, the positive point cloud submap $P = \{(x_i, y_i, z_i)\}_{i=1}^M \in \mathbb{R}^{M \times 3}$ where M is the number of 3D points, the set of negative point cloud submaps $N = \{N_i\}_{i=1}^N$ where $N_i \in \mathbb{R}^{M \times 3}$, the camera pose $y_{gt} = (\mathbf{R}, \mathbf{t}) \in SE(3)$, and the camera intrinsic matrix $\mathbf{K} \in \mathbb{R}^{3 \times 3}$. The positive submap indicates the one that contains the point clouds of the scene in the image, while the negative submaps are randomly taken from other regions. Figure 1 describes the overall pipeline of our approach during inference, and Figure 2 illustrates our network architecture for end-to-end training. We below explain its details in the training and inference process. The implementation details are presented in Section 4.3.

3.1. Feature extraction

Since we follow the coarse-to-fine matching scheme, we extract both local and global descriptors from both image and point cloud submaps. Our feature extraction network for images and point clouds are respectively based on the Swin Transformer [30] and the Fast Point Transformer [37]. We use Transformer-based extractors [52], since they are known to be robust to modality differences, making it possible to extract features that share the embedding space regardless of input modality [40,55]. Then, the global descriptors are obtained by aggregation of local descriptors. These global descriptors are used to reduce search space from the 3D global reference map by retrieving relevant submaps from the database, which we call *global matching*.

2D image descriptors. We first divide the input image I into tokens of the size of 4×4 . For feature extraction of each 2D patch, we use the modified Swin-S Transformer with a window size of 32, which consists of 4-stage Swin-T

Transformer blocks and an FC layer. The output resolution is then $\frac{H}{g} \times \frac{W}{g}$, where g is our 2D patch size, and we get a descriptor of 128-dim for each 2D patch. The global descriptor is obtained by applying a NetVLAD [2] layer to the FC layer. Finally, we obtain a 256-dim global descriptor G_I and 128-dim 2D patch descriptors.

Removal of invisible 3D points. Before extracting point cloud features, we exclude the parts of point cloud P that are invisible in the query image I . In order to identify occluded points, one intuitive idea is to keep only the closest 3D points that can be projected per image pixel to the depth map $D \in \mathbb{R}^{H \times W \times 1}$, which is constructed by projecting the 3D point cloud P to the image frame. However, due to much empty space in a 3D point cloud, points from invisible objects can be projected onto other pixels and blended together, making it difficult to identify occluded points. To effectively deal with this problem, we use min-pooling and max-pooling together. Our *Invisible 3D Point Removal (IPR)* algorithm removes the points where the depth value does not change when applying min-pooling and max-pooling sequentially to the depth map as follows:

$$\text{IPR}(P) = P[\text{maxpool}(\text{minpool}(D)) = D]. \quad (1)$$

We set the kernel size of $s = 9$ for both poolings. For each 3D point in the point cloud P , called a query point, the first min-pooling searches for the closer 3D points. Then the max-pooling checks whether those closer 3D points surround the query point. This prevents 3D points from invisible objects from being projected onto other pixels and blended together due to the large amount of empty space in the 3D point cloud. The time complexity of the IPR algorithm is $O(M + HWs^2)$, where $O(M)$ is for projecting the point cloud, $O(HWs^2)$ for building a depth map and pooling, and $O(M)$ for removing points. Figure 3 shows some examples of invisible point removal. The efficiency and effectiveness of our IPR algorithm can be seen in Table 4 and discussed in Section 4.6.

3D point cloud descriptors. We voxelize the point cloud submaps using Minkowski Engine [13] to form the input tokens. After applying the IPR algorithm, we extract 3D local features using the Fast Point Transformer. We add two branches at the end of this network; one is followed by a NetVLAD layer that aggregates 3D point descriptors into a global descriptor, and the other is for 2D patch classification, which will be described in Section 3.2. Finally, we obtain 128-dim local descriptors per 3D point and a 256-dim global descriptor of the point cloud submap.

3.2. 2D patch classification

2D patch classification determines whether each 3D point is included in the image or not, and if yes, to which 2D patch in the image it belongs. It can significantly reduce the search space to the patch-level during inference. The

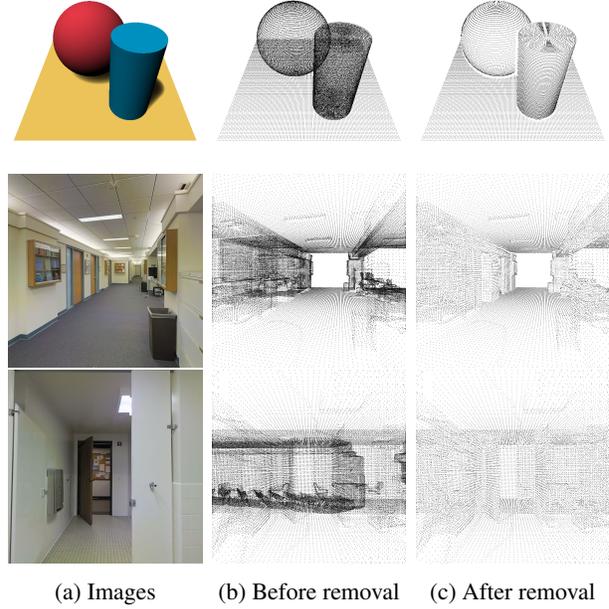


Figure 3: Examples of IPR results. (a) Input images with one toy example and one 2D-3D-S example [3]. (b) Corresponding point cloud submaps. (c) Submaps after removing invisible points from (b).

number of classes to classify is the number of 2D patches plus one class indicating the 3D point does not belong to the image. After concatenating the 3D point descriptors and the 2D patch descriptors, we apply an MLP with softmax to derive classification scores, as shown in Figure 2.

3.3. 2D pixel features with positional encoding

In order to increase the number of inliers during matching the features between 2D pixels and 3D points, it is desired to have dense features in at least one modality. Intuitively, the 2D image is better than the 3D point cloud for dense feature extraction, since the image is sampled based on a dense grid of pixels. Assuming that the 2D patch-level features in the previous section are expressive enough to represent 2D pixel-level features, we design a *2D patch-to-2D pixel function* that returns the 2D pixel features from the corresponding 2D patch feature using positional encoding. We use a simple two-layer MLP, pixel-MLP in Figure 2, whose inputs are the 2D patch features F_{patch} and the positional encoding $PE(r(\mathbf{x}))$ in 2D pixel coordinates \mathbf{x} . r is a function that converts \mathbf{x} to coordinates within the patch containing \mathbf{x} , normalized to $[-1, 1]$. We adopt the multi-dimensional positional encoding used in NeRF [34] with the frequency range of $L = 6$:

$$F_{\mathbf{x}} = \text{MLP}(F_{\text{patch}}, PE(r(\mathbf{x}))), \quad (2)$$

$$PE(\mathbf{p}) = \{\sin(2^f \pi \mathbf{p}), \cos(2^f \pi \mathbf{p}) \mid f \in [0, L - 1]\}, \quad (3)$$

where $\mathbf{p} = r(\mathbf{x})$. Since there are multiple frequencies in the positional encoding, 2D pixel features are sufficiently distinctive from 2D patch features. To increase the generalizability, we use the coordinates obtained by projecting 3D points onto the image directly. It can learn features regardless of the pixel grid, which makes features expressed smoothly with continuous positional encoding. However, for efficiency, we only extract features for each pixel coordinate and compare them with 3D point features at test.

Our approach enables hierarchical matching; we first compare 3D point features with 2D patch features, and then with 2D pixel features within the 2D patch, to efficiently perform dense local matching. As a result, we can reduce the search space from $O(HW)$ to $O(\frac{HW}{g^2} + g^2)$ for a patch size g , and eventually to $O(\sqrt{HW})$ if $g = (HW)^{\frac{1}{4}}$.

3.4. A differentiable PnP layer

Finally, a differentiable PnP solver learns weights w_i to find the best 2D-3D correspondences for high-quality 6-DoF pose estimation. Specifically, the PnP solver is to find the optimal camera rotation matrix \mathbf{R}^* and translation \mathbf{t}^* that minimizes the sum of squared weighted reprojection errors of the 2D-3D correspondences $(\mathbf{x}_i, \mathbf{X}_i)$.

$$y^* = (\mathbf{R}^*, \mathbf{t}^*) = \underset{y=(\mathbf{R}, \mathbf{t}) \in SE(3)}{\operatorname{argmin}} \frac{1}{2} \sum_i \underbrace{\|w_i \circ (\pi(\mathbf{R}\mathbf{X}_i + \mathbf{t}) - \mathbf{x}_i)\|^2}_{f_i(y) \in \mathbb{R}^2}, \quad (4)$$

where \circ is the element-wise product, π is a function that projects 3D points of the camera coordinate system to 2D points of the image coordinate system based on the camera intrinsics \mathbf{K} , and $f_i(y)$ is the weighted reprojection error. To this end, we adopt EPro-PnP [11] for end-to-end training.

3.5. Training

Our loss function is defined by four terms of global matching loss, patch and pixel losses, and a PnP loss:

$$L_{Total} = \alpha \mathcal{L}_{Global} + \beta \mathcal{L}_{Patch} + \gamma \mathcal{L}_{Pixel} + \mathcal{L}_{KL}. \quad (5)$$

We set $\alpha = \beta = \gamma = 0.5$ for experiments.

For global feature matching, we adopt a *Triplet margin loss* [48] defined as

$$\mathcal{L}_{Triplet}(a, p, n) = [m + D(a, p) - D(a, n)]_+, \quad (6)$$

where $D(\cdot, \cdot)$ is the Euclidean distance, $[x]_+ = \max(x, 0)$ is the hinge loss, and m is a hyper-parameter for the margin (we set $m = 0.4$). a , p , and n indicate anchor, positive, and negative, respectively. We construct a set of training triplet tuples $\{I, P, N\}$, where I is an image, and P and N respectively represent its positive point cloud submap and a set of negative submaps, for which we randomly sample

submaps from the regions outside of P . Then, our global feature matching loss is defined as

$$\mathcal{L}_{Global} = \mathcal{L}_{Triplet}(G_I, G_P, G_N), \quad (7)$$

where G_I , G_P , and G_N are global features extracted from I , P , and N , respectively.

For the patch loss \mathcal{L}_{Patch} , we use the cross-entropy loss since the task is a 2D patch classification.

2D pixel feature learning also uses triplet margin loss. For a 3D point feature F_{point} , $F_{\mathbf{x}}$ and $F_{\mathbf{x}_n}$ are respectively positive and negative features, and the loss is

$$\mathcal{L}_{Pixel} = \mathcal{L}_{Triplet}(F_{point}, F_{\mathbf{x}}, F_{\mathbf{x}_n}). \quad (8)$$

A negative pixel \mathbf{x}_n is randomly sampled from the pixels that are at least $\frac{g}{2}$, half the patch size away from \mathbf{x} .

For training of the differentiable PnP solver, we view the PnP output as a pose distribution. Letting y as the estimated pose and y_{gt} as the ground-truth pose: the following KL divergence is minimized as a training loss.

$$\mathcal{L}_{KL} = \frac{1}{2} \sum_{i=1} \|f_i(y_{gt})\|^2 + \log \int \exp \left(-\frac{1}{2} \sum_{i=1} \|f_i(y)\|^2 \right) dy. \quad (9)$$

3.6. Inference

At inference, we predict the 6-DoF camera coordinates of a query image using the point cloud submaps as a database. We first extract 2D patch-level local features and a global descriptor as presented in 3.1. We retrieve the top 4 3D point cloud submaps from the database using the global descriptor. The database is constructed by combining all of the training data from test regions that are not used during training, and the global descriptor and the 3D point descriptors are extracted from the point cloud submap with invisible points removed for the associated training image. Then, we perform 2D patch classification of each 3D point in the retrieved point cloud submaps to determine whether the 3D points of the candidates belong to which grid of the image, and calculate the 2D pixel coordinates of 3D points that are determined to belong to the image using positional encoding. We put all 2D-3D correspondences into the differentiable PnP layer in Section 3.4, and finally predict the 6-DoF camera pose of the query image. Figure 1 illustrates the overall pipeline of our approach.

4. Experiments

4.1. Benchmark datasets

Existing benchmarks for image-based visual localization [9, 18, 19, 32] are often collected at a few different times at the same location. This may lead to positional errors in the pose sensors and poor pose alignment between sequences,

Dataset	Fold	Training		Test	
		Areas	# of data	Areas	# of data
2D-3D-S [3]	1	1, 2, 3, 4, 6	8,134	5	6,614
	2	1, 3, 5, 6	6,482	2, 4	10,123
	3	2, 4, 5	7,218	1, 3, 6	8,511
KITTI [19]	-	00 - 08	7,950	09, 10	1,052

Table 1: The dataset configuration of 2D-3D-S [3] and KITTI [19].

resulting in inaccurate 2D-3D correspondences. Therefore, we curate new benchmark datasets for the 2D image to 3D point cloud localization task.

Our benchmark datasets are based on the indoor Stanford 2D-3D-Semantic (2D-3D-S) [3] and the outdoor KITTI [19] dataset, as shown in Figure 4. The 2D-3D-S dataset contains 6 large-scale indoor areas scanned by the MatterPort camera, and each area includes the global point cloud map provided in the Stanford large-scale 3D Indoor Spaces dataset (S3DIS) [4]. We sample RGB images for training and testing so that the angles between the images taken at the same location do not differ by less than 30° and 15° , respectively, following the InLoc [51] setting. For the coarse-to-fine matching approach, we provide the global point cloud maps along with small submaps. We divide the global point cloud map into fixed-size submaps visible in the training image frames within 10m and downsample to have $M = 65, 536$ points. We follow the 3-fold cross-validation scheme of 2D-3D-S as shown in Table 1, which evaluates the performance of pose estimation for unseen places during training. In each fold, the database for the test areas is constructed by collecting the training data for those areas that are not used during training. For example, the database for the Fold 3 consists of the training data from Areas 2, 4, and 5 combined.

The KITTI dataset captures large-scale outdoor environments by driving on the road in a medium-sized city of rural area. There are 11 sequences numbered 00 to 10 that have ground-truth poses, and we split the sequences 00 to 08 for training and the sequences 09 to 10 for testing. Images are collected every 2m from the trajectory with no preprocessing. We construct a point cloud submap with the global coordinates by accumulating 3D LiDAR data and sampling the points visible in the image frame within $30m$, and down-sampling each submap to $M = 65, 536$ points. More details are described in the supplementary material.

4.2. Metrics

We evaluate two types of performances: the pose estimation of a query image (i) within the entire point cloud (*i.e.* visual localization) and (ii) within a single point cloud (*i.e.* image-to-point cloud registration). For the visual localization task, we report the percentage of query images that are correctly localized within three thresholds of trans-

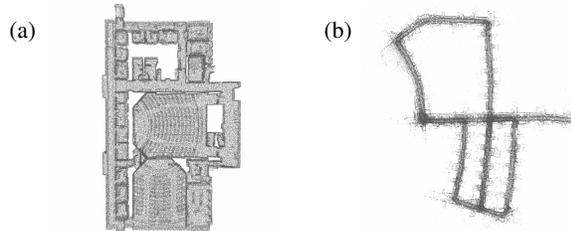


Figure 4: Global point cloud maps of (a) 2D-3D-S [3] (Area 2) and (b) KITTI [19] (Sequence 05).

lation and rotation: (0.1m, 1.0°), (0.25m, 2.0°), and (1.0m, 5.0°). This performance measurement follows the hardest evaluation method for the indoor environment of the *Long-Term Visual Localization*¹ challenge. For the image-to-point cloud registration task, we report average Relative Translational Error (RTE (m)) and average Relative Rotation Error (RRE ($^\circ$)).

4.3. Implementation details

During training, we are given an image, the positive point cloud submap, a set of negative point cloud submaps, the camera pose, and the camera intrinsic matrix. To shorten the training time, we organize a batch by sampling training images with the positive point cloud submap that can be used as negatives for each other. For 2D patch classification and 2D pixel feature learning, we generate ground-truth by projecting 3D points from the positive point cloud submap onto the image using the camera intrinsic matrix and the relative pose between the image and the point cloud. If the 3D point is projected within the image frame, we specify the 2D patch number as a class number starting at 1 and set it to the ground-truth along with the 2D pixel coordinates. Otherwise, we set the 2D patch number of the 3D point to the 0-th class, indicating that it does not belong to the image frame. During training, these 2D patch numbers are used for 2D patch classification and the 2D pixel coordinates are used for triple tuple setup. The loss function is minimized using Adam optimizer [26], for which the initial learning rate is 10^{-4} and divided by 10 at the epoch given in the LR scheduler steps.

4.4. Results of visual localization

We compare our approach with six image-to-image localization methods, one image-to-RGB-D method [23] and one RGB-D-to-RGB-D method [7]. The image-to-image baselines consist of a DSAC-based DSAC* [7], SfM methods such as COLMAP [47] and Kapture [22], and the state-of-the-art local feature extraction methods including SuperPoint [14] + NetVLAD [2], D2-Net [15] + NetVLAD, and HLoc [42, 43]. The SfM baselines construct a 3D reference

¹<https://www.visuallocalization.net/benchmark/>

Model	Fold 1			Fold 2			Fold 3			Runtime (sec)
	(0.1, 1.0)	(0.25, 2.0)	(1.0, 5.0)	(0.1, 1.0)	(0.25, 2.0)	(1.0, 5.0)	(0.1, 1.0)	(0.25, 2.0)	(1.0, 5.0)	
COLMAP [47]	68.78	73.21	75.40	79.75	82.62	84.21	65.39	70.36	73.43	9.377
Kapture [22]	72.80	77.49	79.56	81.39	84.88	86.58	66.96	72.28	74.53	1.911
SuperPoint [14] + NetVLAD [2]	81.49	83.26	84.23	88.57	89.89	90.37	78.86	81.20	82.25	3.892
HLoc [42, 43]	84.14	86.50	87.86	90.47	92.10	92.86	80.84	83.80	85.18	6.932
D2-Net [†] [15] + NetVLAD	80.06	82.67	84.00	86.81	88.88	89.89	76.13	79.34	80.99	5.815
DSAC* _{image-to-image} [7]	0.00	0.00	0.29	0.16	1.56	11.41	0.74	5.18	19.84	0.738
PCLoc [†] [23]	76.34	80.99	83.87	76.97	80.19	83.92	74.26	77.97	80.41	14.27
DSAC* _{RGB-D-to-RGB-D} [†]	0.00	0.00	0.26	0.00	0.52	8.94	0.86	3.75	13.54	5.874
EP2P-Loc (Ours)	87.59	91.65	92.89	92.47	93.88	94.38	85.24	87.38	89.43	3.482

Table 2: Experimental results of visual localization accuracy and runtime on the 2D-3D-S dataset [3]. Please refer to Table 1 for the folds. The threshold units are (m, \circ) . [†] indicates models that use depth maps during training.

Model	RTE (m)	RRE (\circ)
Direct Regression	4.94 ± 2.87	21.98 ± 31.97
Monodepth2 [21] + USIP [28]	30.4 ± 42.9	140.6 ± 157.8
Monodepth2 + GT-ICP [6, 12]	2.9 ± 2.5	12.4 ± 10.3
2D3D-MatchNet [16] [†]	752.5 ± 6053.3	117.9 ± 52.1
DeepI2P [29]	3.28 ± 3.09	7.56 ± 7.63
EP2P-Loc (Ours)	1.32 ± 1.13	4.11 ± 5.46

Table 3: Experimental results of image-to-point cloud registration on the KITTI dataset [19]. [†] indicates the results of our implementation.

map using COLMAP on the database images. The pipeline of local feature extraction methods consists of two steps: (i) retrieving the top 20 nearest database images through NetVLAD and (ii) estimating pose by applying the local feature matching results to the PnP solver [17, 27].

Table 2 summarizes the results. Our method outperforms existing image-to-image, image-to-RGB-D, or RGB-D-to-RGB-D localization approaches in all three-fold configurations with all three threshold settings. Since the coordinate regression errors of the DSAC* are proportional to the area size, it becomes unreliable in large-scale localization. In addition, the RGB-D-to-RGB-D version of DSAC* also utilizes depth information of input images, but scores lower than the image-to-image version, since it highly downsamples the given depth map for efficiency but loses fine structure. The estimated pose of PCLoc [23] depends on the top-1 candidate pose selected through the pose correction and pose verification, and does not fully utilize the given image, resulting in a low inlier problem. Another observation is that the image-to-image localization methods show better accuracy than DSAC* and PCLoc. They use the SfM model to construct 2D-3D correspondences and are guided by SIFT [31] matching to triangulation.

4.5. Results of image-to-point cloud registration

This task takes a query image and a single point cloud as inputs and estimates the relative pose between the image

and the point cloud. As baselines, we first select two state-of-the-art methods for 2D-3D registration: 2D3D-MatchNet [16] and DeepI2P [29]. Since the code of 2D3D-MatchNet is not released, we implement it by ourselves. We also add three baselines of DeepI2P: (i) Direct Regression, (ii) Monodepth2 [21] + USIP [28], and (iii) Monodepth2 + GT-ICP. Monodepth2 is the state-of-the-art depth estimation method, and USIP is the method to estimate the pose between the depth map and the point cloud. GT-ICP indicates the Interactive Closest Point (ICP) method [6, 12] that uses the ground-truth (GT) relative pose for initialization. These three baselines have unfair advantages over the other methods by using additional data or GTs. Please refer to [29] for more details of the baselines.

As shown in Table 3, our method shows the state-of-the-art performance on the KITTI dataset [19]. 2D3D-MatchNet shows the worst performance; as pointed out in [29], 2D3D-MatchNet suffers from a low inlier rate in the correspondences, since it should learn to overcome the drastic discrepancy between image-based SIFT [31] and point cloud-based ISS [56] keypoints. Moreover, the inlier ratio could further decrease in the KITTI dataset, which mostly consists of front-view images where similar patterns repeat and far buildings are shown relatively small. This requires extremely many iterations of PnP and RANSAC [17], and thus its localization is unstable. The superiority of our method over DeepI2P shows that our idea of explicitly calculating 3D points to 2D pixel coordinates can perform better than DeepI2P’s grid classification of each 3D point.

4.6. Ablation study

IPR algorithm. In the Table 4, we show the quantitative result of our IPR algorithm on the 2D-3D-S dataset [3]. Our IPR algorithm removes 68.9% of the points within 10.6 ms per submap during preprocessing. Even though the number of 3D points is reduced by a third, the number of inliers is reduced by only 12.5%, the performance is similar or better, and the runtime is reduced by 50.12%.

Differentiable PnP solver. To show the effect of using a

Model	Fold 1			Fold 2			Fold 3			# of points	# of inliers	Runtime (sec)
	(0.1, 1.0)	(0.25, 2.0)	(1.0, 5.0)	(0.1, 1.0)	(0.25, 2.0)	(1.0, 5.0)	(0.1, 1.0)	(0.25, 2.0)	(1.0, 5.0)			
EP2P-Loc w/o IPR	88.12	90.33	92.75	91.57	94.12	94.89	83.23	86.36	88.29	65.5K	3.12K	6.981
EP2P-Loc w/ EPnP [27]	85.34	87.51	90.23	90.23	92.34	93.78	83.56	84.96	87.13	20.4K	0.32K	12.53
EP2P-Loc (Ours)	87.59	91.65	92.89	92.47	93.88	94.38	85.24	87.38	89.43	20.4K	2.73K	3.482

Table 4: Ablation study on the 2D-3D-S dataset [3]. The threshold units are (m, \circ) . The number of points is reported with the average value per point cloud submap in the training set. The number of inliers and runtime are reported with the average value per image in the test set, and runtime is measured with the overall pipeline including data loading to pose estimation.

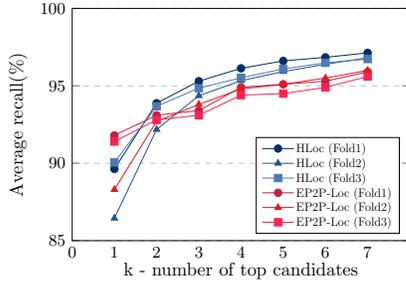


Figure 5: Comparison of average recalls for global matching on the 2D-3D-S dataset [3].

differentiable PnP solver in this task, we experiment with an EPnP solver [27] with RANSAC [17]. As shown in Table 4, the differentiable PnP solver outperforms an EPnP solver with the same 2D-3D correspondences, which weights the correspondences and efficiently estimating poses with a high number of inliers.

4.7. Analysis

Number of candidates to retrieval in global matching. We analyze the performance of the global descriptor with a different number of candidates. We evaluate whether the RTE (m) for the top k most similar candidates is lower than $0.1m$, and plot the recall curves of each fold of the 2D-3D-S dataset [3], as shown in Figure 5. To show the effectiveness of our global descriptor, we compare our results with HLoc [42, 43], which is based on the coarse-to-fine approach as ours. Our method performs comparably to image-to-image matching, and outperforms HLoc, especially for the top 1. Moreover, our method can estimate a more accurate pose, as shown in Table 2. While the recall value improves as k increases, we construct the pipeline using $k = 4$, considering the trade-off.

2D patch classification and 2D pixel coordinate calculation. Table 5 shows the accuracy of the 2D patch classification and the 2D pixel coordinate calculation on the 2D-3D-S and KITTI [19] data. The 2D-3D-S dataset, which is based on RGB-D data, has a more sophisticated data distribution than the KITTI dataset, which can lead to more accurate camera pose estimation.

Runtime comparison. Table 2 presents that our model

Dataset	Fold	2D patch classification	2D pixel calculation
2D-3D-S [3]	1	78.2	41.5
	2	82.3	43.7
	3	70.8	38.5
KITTI [19]	-	63.8	34.9

Table 5: Results of 2D patch classification and 2D pixel calculation on the 2D-3D-S [3] and the KITTI [19] datasets.

achieves superior performance while showing fast runtime compared to other state-of-the-art methods. DSAC* [7] is a regression-based method, and Kapture [22] uses GeM pooling [39] instead of NetVLAD [2], resulting in faster runtime but showing inferior performance. HLoc attains lower performance than our model but takes approximately twice the runtime.

5. Conclusion

To estimate the camera pose using a PnP solver for visual localization, it is necessary to obtain the 2D pixels of an image and the 3D point correspondences of a 3D reference map as input. In this work, we propose a novel approach to solve the large-scale visual localization task by mitigating representational differences between 2D and 3D before feature extraction and finding all 2D-3D correspondences to increase the number of inliers, reducing memory usage and computational cost. And we adopt an end-to-end trainable PnP solver, for the first time in this task, to learn to select good 2D-3D pairs for pose estimation utilizing the ground-truth of 6-DoF camera pose during training. We experiment on our benchmark datasets based on large-scale indoor and outdoor environments, and show that our method achieves the state-of-the-art performance compared to other previous visual localization and image-to-point cloud methods.

Acknowledgement. We sincerely thank Junhyug Noh, Dongyeon Woo, Minui Hong, and Hyungsuk Lim for their constructive comments. This work was supported by National Research Foundation of Korea (NRF) grant (No.2023R1A2C2005573) and Institute of Information & Communications Technology Planning & Evaluation (IITP) grant (No.2019-0-01082, No.2019-0-01309, No.2021-0-01343) funded by the Korea government (MSIT). Gunhee Kim is the corresponding author.

References

- [1] Sameer Agarwal, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M Seitz, and Richard Szeliski. Building rome in a day. *Communications of the ACM*, 54(10):105–112, 2011.
- [2] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Netvlad: Cnn architecture for weakly supervised place recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5297–5307, 2016.
- [3] I. Armeni, A. Sax, A. R. Zamir, and S. Savarese. Joint 2D-3D-Semantic Data for Indoor Scene Understanding. *ArXiv e-prints*, Feb. 2017.
- [4] Iro Armeni, Ozan Sener, Amir R Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3d semantic parsing of large-scale indoor spaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1534–1543, 2016.
- [5] Thuan Bui Bach, Tuan Tran Dinh, and Joo-Ho Lee. Featloc: Absolute pose regressor for indoor 2d sparse features with simplistic view synthesizing. *ISPRS Journal of Photogrammetry and Remote Sensing*, 189:50–62, 2022.
- [6] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor fusion IV: control paradigms and data structures*, volume 1611, pages 586–606. International Society for Optics and Photonics, 1992.
- [7] Eric Brachmann and Carsten Rother. Visual camera relocalization from rgb and rgb-d images using dsac. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):5847–5865, 2022.
- [8] Samarth Brahmabhatt, Jinwei Gu, Kihwan Kim, James Hays, and Jan Kautz. Geometry-aware learning of maps for camera localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2616–2625, 2018.
- [9] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nusenes: A multimodal dataset for autonomous driving. *arXiv preprint arXiv:1903.11027*, 2019.
- [10] Daniele Cattaneo, Matteo Vaghi, Simone Fontana, Augusto Luis Ballardini, and Domenico Giorgio Sorrenti. Global visual localization in lidar-maps through shared 2d-3d embedding space. In *International Conference on Robotics and Automation (ICRA)*, pages 4365–4371, 2020.
- [11] Hansheng Chen, Pichao Wang, Fan Wang, Wei Tian, Lu Xiong, and Hao Li. Epro-pnp: Generalized end-to-end probabilistic perspective-n-points for monocular object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [12] Yang Chen and Gérard Medioni. Object modeling by registration of multiple range images. *Image and Vision Computing*, 10(3):145–155, 1992.
- [13] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3075–3084, 2019.
- [14] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 224–236, 2018.
- [15] Mihai Dusmanu, Ignacio Rocco, Tomas Pajdla, Marc Pollefeys, Josef Sivic, Akihiko Torii, and Torsten Sattler. D2-net: A trainable cnn for joint detection and description of local features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [16] Mengdan Feng, Sixing Hu, Marcelo H Ang, and Gim Hee Lee. 2d3d-matchnet: Learning to match keypoints across 2d image and 3d point cloud. In *International Conference on Robotics and Automation (ICRA)*, pages 4790–4796. IEEE, 2019.
- [17] Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, june 1981.
- [18] Whye Kit Fong, Rohit Mohan, Juana Valeria Hurtado, Lubing Zhou, Holger Caesar, Oscar Beijbom, and Abhinav Valada. Panoptic nusenes: A large-scale benchmark for lidar panoptic segmentation and tracking. *arXiv preprint arXiv:2109.03805*, 2021.
- [19] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [20] Ben Glocker, Shahram Izadi, Jamie Shotton, and Antonio Criminisi. Real-time rgb-d camera relocalization. In *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 173–179. IEEE, 2013.
- [21] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3828–3838, 2019.
- [22] Martin Humenberger, Yohann Cabon, Nicolas Guerin, Julien Morat, Jérôme Revaud, Philippe Rerole, Noé Pion, Cesar de Souza, Vincent Leroy, and Gabriela Csurka. Robust image retrieval-based visual localization using kapture. *arXiv preprint arXiv:2007.13867*, 2020.
- [23] Janghun Hyeon, Joohyung Kim, and Nakju Doh. Pose correction for highly accurate visual localization in large-scale indoor spaces. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15974–15983, October 2021.
- [24] Heinly Jared, Johannes L Schonberger, Enrique Dunn, and Jan-Michael Frahm. Reconstructing the world in six days. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [25] Alex Kendall and Roberto Cipolla. Geometric loss functions for camera pose regression with deep learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5974–5983, 2017.

- [26] Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- [27] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Epnp: An accurate o(n) solution to the pnp problem. *International Journal of Computer Vision (IJCV)*, 81(2):155–166, 2009.
- [28] Jiaxin Li and Gim Hee Lee. Usip: Unsupervised stable interest point detection from 3d point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 361–370, 2019.
- [29] Jiaxin Li and Gim Hee Lee. DeepI2P: Image-to-Point Cloud Registration via Deep Classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15960–15969, 2021.
- [30] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (CVPR)*, pages 10012–10022, 2021.
- [31] David G Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision (IJCV)*, 60(2):91–110, 2004.
- [32] Will Maddern, Geoff Pascoe, Chris Linegar, and Paul Newman. 1 Year, 1000km: The Oxford RobotCar Dataset. *The International Journal of Robotics Research (IJRR)*, 36(1):3–15, 2017.
- [33] Iaroslav Melekhov, Juha Ylioinas, Juho Kannala, and Esa Rahtu. Image-based localization using hourglass networks. In *Proceedings of the IEEE International Conference on Computer Vision Workshops (ICCVW)*, pages 879–886, 2017.
- [34] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [35] Hyeonwoo Noh, Andre Araujo, Jack Sim, Tobias Weyand, and Bohyung Han. Large-scale image retrieval with attentive deep local features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3456–3465, 2017.
- [36] Onur Ozyesil, Vladislav Voroninski, Ronen Basri, and Amit Singer. A survey of structure from motion. *arXiv preprint arXiv:1701.08493*, 2017.
- [37] Chunghyun Park, Yoonwoo Jeong, Minsu Cho, and Jaesik Park. Fast point transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16949–16958, 2022.
- [38] Quang-Hieu Pham, Mikaela Angelina Uy, Binh-Son Hua, Duc Thanh Nguyen, Gemma Roig, and Sai-Kit Yeung. Lcd: Learned cross-domain descriptors for 2d-3d matching. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 34, pages 11856–11864, 2020.
- [39] F. Radenović, G. Toliás, and O. Chum. Fine-tuning cnn image retrieval with no human annotation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(7):1655–1668, 2019.
- [40] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021.
- [41] Jerome Revaud, Philippe Weinzaepfel, César Roberto de Souza, and Martin Humenberger. R2D2: repeatable and reliable detector and descriptor. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [42] Paul-Edouard Sarlin, Cesar Cadena, Roland Siegwart, and Marcin Dymczyk. From coarse to fine: Robust hierarchical localization at large scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12716–12725, 2019.
- [43] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning feature matching with graph neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [44] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Fast image-based localization using direct 2d-to-3d matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 667–674. IEEE, 2011.
- [45] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Efficient & effective prioritized matching for large-scale image-based localization. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 39(9):1744–1756, 2016.
- [46] Torsten Sattler, Qunjie Zhou, Marc Pollefeys, and Laura Leal-Taixe. Understanding the limitations of cnn-based absolute camera pose regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3302–3312, 2019.
- [47] Johannes L. Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [48] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–823, 2015.
- [49] Jamie Shotton, Ben Glocker, Christopher Zach, Shahram Izadi, Antonio Criminisi, and Andrew Fitzgibbon. Scene coordinate regression forests for camera relocalization in rgb-d images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2930–2937, 2013.
- [50] Noah Snavely, Steven M Seitz, and Richard Szeliski. Modeling the world from internet photo collections. *International Journal of Computer Vision (IJCV)*, 80(2):189–210, 2008.
- [51] Hajime Taira, Masatoshi Okutomi, Torsten Sattler, Mircea Cimpoi, Marc Pollefeys, Josef Sivic, Tomas Pajdla, and Akihiko Torii. Inloc: Indoor visual localization with dense matching and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7199–7209, 2018.
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Il-

- lia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems (NIPS)*, volume 30. Curran Associates, Inc., 2017.
- [53] Bing Wang, Changhao Chen, Zhaopeng Cui, Jie Qin, Chris Xiaoxuan Lu, Zhengdi Yu, Peijun Zhao, Zhen Dong, Fan Zhu, Niki Trigoni, et al. P2-net: Joint description and detection of local features for pixel and point matching. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [54] Xiaoxia Xing, Yinghao Cai, Tao Lu, Shaojun Cai, Yiping Yang, and Dayong Wen. 3dtnet: Learning local features using 2d and 3d cues. In *International Conference on 3D Vision (3DV)*, pages 435–443, 2018.
- [55] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Bin Xiao, Ce Liu, Lu Yuan, and Jianfeng Gao. Unified contrastive learning in image-text-label space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19163–19173, 2022.
- [56] Yu Zhong. Intrinsic shape signatures: A shape descriptor for 3d object recognition. In *Proceedings of the IEEE International Conference on Computer Vision Workshops (ICCVW)*, pages 689–696. IEEE, 2009.