

# 360Loc: A Dataset and Benchmark for Omnidirectional Visual Localization with Cross-device Queries

Huajian Huang<sup>1\*</sup> Changkun Liu<sup>1\*</sup> Yipeng Zhu<sup>1</sup> Hui Cheng<sup>2</sup> Tristan Braud<sup>1</sup> Sai-Kit Yeung<sup>1</sup>  
<sup>1</sup>The Hong Kong University of Science and Technology <sup>2</sup>Sun Yat-sen University  
\* equal contribution

{hhuangbg, cliudg, yzhudg}@connect.ust.hk, chengh9@mail.sysu.edu.cn, {braudt, saikit}@ust.hk

## Abstract

Portable 360° cameras are becoming a cheap and efficient tool to establish large visual databases. By capturing omnidirectional views of a scene, these cameras could expedite building environment models that are essential for visual localization. However, such an advantage is often overlooked due to the lack of valuable datasets. This paper introduces a new benchmark dataset, 360Loc, composed of 360° images with ground truth poses for visual localization. We present a practical implementation of 360° mapping combining 360° images with lidar data to generate the ground truth 6DoF poses. 360Loc is the first dataset and benchmark that explores the challenge of cross-device visual positioning, involving 360° reference frames, and query frames from pinhole, ultra-wide FoV fisheye, and 360° cameras. We propose a virtual camera approach to generate lower-FoV query frames from 360° images, which ensures a fair comparison of performance among different query types in visual localization tasks. We also extend this virtual camera approach to feature matching-based and pose regression-based methods to alleviate the performance loss caused by the cross-device domain gap, and evaluate its effectiveness against state-of-the-art baselines. We demonstrate that omnidirectional visual localization is more robust in challenging large-scale scenes with symmetries and repetitive structures. These results provide new insights into 360-camera mapping and omnidirectional visual localization with cross-device queries. Project Page and dataset: <https://huajianup.github.io/research/360Loc/>.

## 1. Introduction

Visual localization refers to predicting the 6DoF absolute pose (translation and rotation) of query images in a known scene. Accurate visual localization has wide applications in augmented reality (AR), navigation, and robotics.

Over the last decade, many visual localization methods have been proposed, including feature matching-based approaches [17, 33, 42, 45, 54], scene coordinate regression [5–7] and absolute pose regressors (APRs) [23, 24, 49]. Much of this progress has been driven by the availability of numerous datasets and benchmarks targeting different challenges, as shown in Table 1. However, existing methods and datasets focus on localization and mapping using pinhole images. Although the merits of 360° camera on visual perception have been recognized [22, 60, 62], the application of 360° cameras for visual localization is still under-explored. Recently, SensLoc [61] started to apply 360° cameras to facilitate data collection, but their pipeline cannot perform omnidirectional localization directly from the 360° images.

This paper introduces 360Loc, a new challenging benchmark dataset to facilitate research on omnidirectional visual localization. The dataset contains 360° images captured in diverse campus-scale indoor and outdoor environments, featuring highly symmetrical and repetitive features, as well as interference of dynamic objects. To capture this dataset, we present a practical pipeline using a portable 360-camera-lidar platform to obtain reliable pose estimations of 360° cameras as ground truth. Although 360° cameras present significant advantages for capturing reference data, real-life applications applying visual localization often rely on traditional cameras. Examples include robots equipped with fisheye cameras and phone-based AR applications using the embedded pinhole camera. This raises the problem of cross-device visual localization on image databases captured with 360° cameras. We thus supplement the reference database composed of 360° images with query frames including pinhole, fisheye and 360° cameras.

We introduce the concept of *virtual camera* to generate high-quality lower-FoV images with different camera parameters from 360° images. This enables a fair comparison of performance among queries from different devices in cross-device visual localization. We adapt existing feature-matching-based methods and APRs to support 360° image queries and benchmark these methods for 360-based cross-

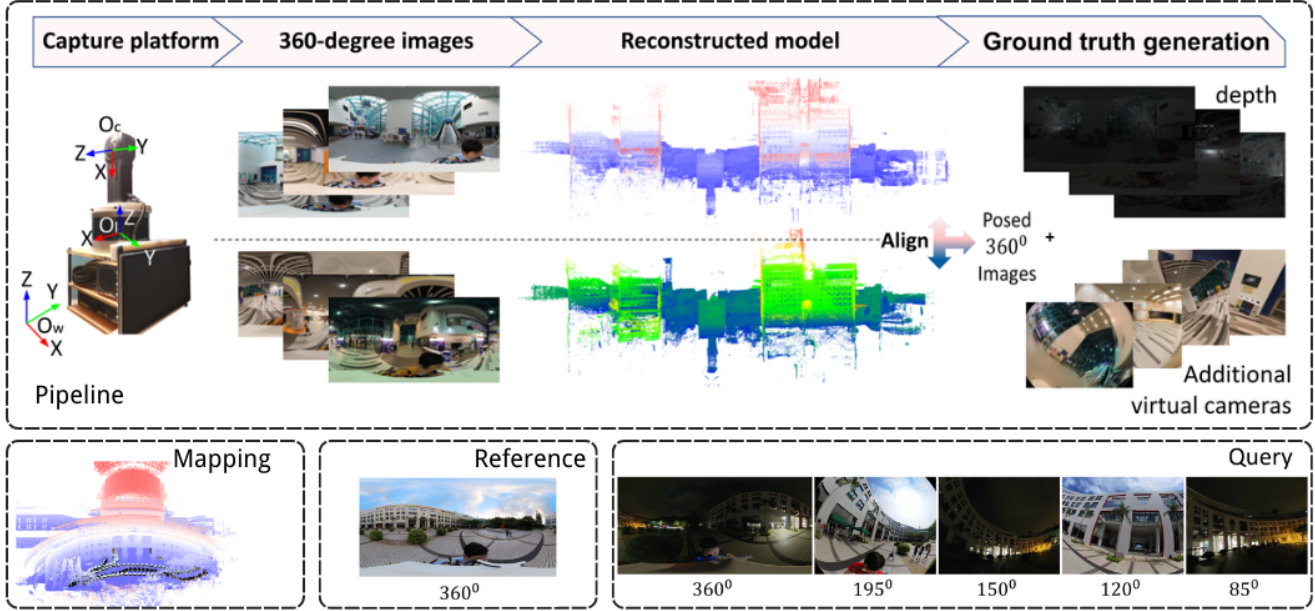


Figure 1. Overview of dataset collection and ground truth generation: 1) Use the platform to collect  $360^\circ$  images and frame-by-frame point clouds. Obtain real-time camera poses; 2) Apply optimization methodology to achieve data registration, resulting in a globally reconstructed point cloud model. Then, align the models in daytime and nighttime to get consistent poses; 3) Perform cropping to get *virtual camera* images and generate corresponding depth images. As a result, 360Loc takes advantage of  $360^\circ$  images for efficient mapping while providing query images in five different camera models in order to analyze the challenge of cross-domain visual localization.

device visual localization. Since different cameras present different imaging patterns, the cross-device domain gap is expected to lead to performance loss. We extend the *virtual camera* approach to data augmentation for end-to-end solutions such as image retrieval (IR) and APRs.

By conducting exhaustive evaluations, we demonstrate the advantages of  $360^\circ$  cameras in reducing ambiguity in visual localization on scenes featuring symmetric or repetitive features. We also show improvements against state-of-the-art (SOTA) baselines using the *virtual camera* method for cross-device visual localization on images databases captured with  $360^\circ$  cameras. These results provide novel insights on mapping using  $360^\circ$  images, enhancing the anti-ambiguity capability of query images, reducing domain gap cross-device in visual localization, and improving the generalization ability of APRs by applying *virtual cameras*.

Our contribution can be summarized as follows:

- We propose a practical implementation of  $360^\circ$  mapping combining lidar data with  $360^\circ$  images for establishing the ground truth 6DoF poses.
- A virtual camera approach to generate high-quality lower-FoV images with different camera parameters from  $360^\circ$  views.
- A novel dataset for cross-device visual localization based on  $360^\circ$  reference images with pinhole, fisheye, and  $360^\circ$  query images.

- Demonstration of our approach’s efficacy over state-of-the-art solutions for visual localization using  $360^\circ$  image databases, resulting in decreased localization ambiguity, reduced cross-device domain gap, and improved generalization ability of APRs.

## 2. Related work

### 2.1. Visual Localization

**Structure-based methods** predict camera poses by establishing 2D-3D correspondences indirectly with local feature extractors and matchers [16, 35, 42, 43, 52, 55] or directly with scene coordinate regression [5–7]. HLoc [42, 43] pipeline scales up to large scenes using image retrieval [1, 3, 18, 20] as an intermediate step, which achieves SOTA accuracy on many benchmarks. This type of approach usually supports pinhole cameras with different intrinsic parameters. However, the performance of  $360^\circ$  and fisheye cameras has not been evaluated before due to the lack of support for  $360^\circ$  cameras in the Structure from Motion (SfM) tools like COLMAP [45] and the lack of datasets for fisheye and  $360^\circ$  cameras. [25–27] are point-cloud-based panorama localization methods for  $360^\circ$  queries but they do not consider cross-device visual localization.

**Absolute Pose Regressors (APRs)** are end-to-end learning-based methods that directly regress the absolute camera

pose from input images without the knowledge of 3D models and establish 2D-3D correspondences. APRs [4, 8, 12, 13, 23, 24, 36, 37, 49, 59] provide faster inference than structure-based methods at the cost of accuracy and robustness [47]. Besides, APRs have generally only been tested on the [9], 7Scenes [50], and Cambridge Landmarks [24] datasets in previous studies. A notable characteristic of these datasets is that the training set and test set images were taken from the same camera. In this paper, we enhance cross-device pose regression for APRs by introducing virtual cameras as a data augmentation technique.

## 2.2. Datasets

The existing dataset has the following limitations. 1). Most datasets [9, 10, 24, 50, 54, 58] do not consider the need for cross-device localization, i.e., query images come from the same camera. Even though some datasets [11, 14, 30, 44, 46, 48, 53, 61] take into account cross-device localization, these devices are only pinhole cameras with different camera intrinsic parameters and do not have particularly large domain-gaps. Compared to [32], our pinhole and fisheye images are extracted from 360° images via virtual cameras, which makes less demands on the device and allows for a fair and more flexible comparison of the effects of different FoVs. In this paper, our 360Loc datasets provide five kinds of queries from pinhole, fisheye and 360° cameras to promote the research of cross-device localization. 2). Now there is no 6DoF visual localization dataset and benchmark considering 360° reference images and 360° query images, even though [2, 25, 38] contain 360° images with 6DoF pose labels, they are not standard visual localization datasets with independent mapping/reference sequences and query sequences like datasets in Table 1. Other datasets [11, 61] use 360° cameras for data collection, in the end they cropped 360° to perspective images and then tailor these images to the classical visual localization pipeline of pinhole cameras. The academic community is mainly driven by benchmarks where all training, reference, and query images are pinhole images because they rely on SfM tools [45] which does not support 360° cameras to obtain ground-truth (GT) and get sparse 3D point cloud models for recovering camera poses. However, we note that the 360° camera can cover the scene with greater efficiency than normal pinhole cameras with narrow Field-of-View (FoV), which makes 360° images particularly suitable as reference images. 3) Although the current dataset has explored the challenges of visual localization from various aspects such as weather variations, day-night transitions, scene changes, and moving individuals and objects [24, 30, 44, 46, 58, 61], there is still insufficient research specifically targeting highly ambiguous environments which contain symmetries, repetitive structures and insufficient textures. Only two indoor datasets [9, 53]

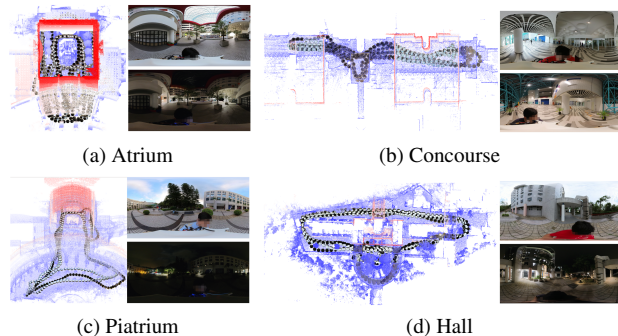


Figure 2. The four scenes in 360Loc, all four scenes contain symmetrical, repetitive structures and moving objects. The camera trajectories are visualized as spheres.

and LaMAR [44] consider challenges in ambiguous environments. In this paper, we studied 4 ambiguous scenes from both indoor and outdoor environments with a scale much larger than dataset [9] (See Figure 2). We conduct exhaustive assessments of image retrieval, local matching localization, and absolute pose regression to show that queries from the 360° camera are harder to obtain plausible solutions than other queries from cameras with narrower FoV.

## 3. The 360Loc Dataset

The 360Loc dataset contains 4 locations from a local university. Figure 2 displays the reference point cloud and example frames from each scene. Atrium is inside a building with a surrounding structure that exhibits a high degree of symmetry and repetition, making it a highly ambiguous environment. Concourse is a large indoor scene with many moving people, which can be used for evaluating the robustness of any localization algorithm in scenes with many moving objects. Piatrium is a scene containing both indoor Atrium and outdoor environments, covering an outdoor piazza with coffee shops, bookstores, and souvenir shops. Hall is a modern building of a student dormitory.

### 3.1. 360 Mapping Platform

We utilized the handheld multimodal data acquisition platform depicted in Figure 1 for data collection. This platform incorporates a 360° camera, a Velodyne VLP-16 multi-line lidar, an NUC mini-computer, and a display screen. Figure 1 also illustrates the relative relationship among the 360° camera coordinate system  $O_c\text{-XYZ}$ , the lidar coordinate system  $O_l\text{-XYZ}$  as well as the world coordinate  $O_w\text{-XYZ}$ . The portable 360 camera equipped on this device can capture high-resolution omnidirectional images with a resolution of 6144 x 3072 (2:1 aspect ratio). It also features a built-in six-axis gyroscope that provides stabilization support, making it suitable for handheld mobile data capture. The Velodyne VLP-16 multi-line lidar has



Dataset	Scale and Environment	Challenges	Reference/Query type	Groundtruth Solution	Accuracy
7Scenes [50]	Small Indoor	None	pinhole / pinhole	RGB-D	$\approx$ cm
RIO10 [58]	Small Indoor	Changes	pinhole / pinhole	VIO	$>$ $d$ m
Baidu Mall [53]	Medium Indoor	People, Ambiguous	pinhole / pinhole	lidar+Manual	$\approx$ $d$ m
Naver Labs [30]	Medium Indoor	People, Changes	pinhole / pinhole	lidar+SfM	$\approx$ $d$ m
InLoc [54]	Medium Indoor	None	pinhole / pinhole	lidar+Manual	$>$ $d$ m
AmbiguousLoc [9]	Small Indoor	Ambiguous	pinhole / pinhole	SLAM	$\approx$ cm
Achen [46]	Large outdoor	People, Day-Night	pinhole / pinhole	SfM	$>$ $d$ m
Cambridge [24]	Medium outdoor	People, Weather	pinhole / pinhole	SfM	$>$ $d$ m
San Francisco [11]	Large outdoor	People, Construction	pinhole / pinhole	SfM+GPS	$\approx$ $m$
NCLT [10]	Medium Outdoor + Indoor	Weather	pinhole / pinhole	GPS+SLAM+lidar	$\approx$ $d$ m
ADVIO [14]	Medium Outdoor+Indoor	People	pinhole / pinhole	VIO+Manual	$\approx$ $m$
ETH3D [48]	Medium Outdoor + Indoor	None	pinhole / pinhole	lidar+Manual	$\approx$ mm
LaMAR [44]	Medium Outdoor+Indoor	People, Weather, Day-Night, Construction, Changes, Ambiguous	pinhole / pinhole	lidar+SfM+VIO	$\approx$ cm
SensLoc [61]	Large Outdoor	People, Weather, Day-Night, Construction, Changes	pinhole / pinhole	SL+VIO+RTK+Gravity	$<$ dm
<b>360Loc (ours)</b>	Medium Outdoor+Indoor	People, Weather, Day-Night, Construction, Changes, <b>Ambiguous</b>	<b>360 / (360 + pinhole + fisheye)</b>	lidar+VIO	$\approx$ cm

Table 1. Overview of popular visual localization datasets. No dataset, besides ours, consider 360° images as reference and query frames from pinhole, ultra-wide FoV fisheye, and 360° cameras.

Symbol	Name	Field of View	Resolution	Type
$c_0$	360	360°	6144 × 3072	reference/query
$c_1$	fisheye1	120°	1280 × 1024	query
$c_2$	fisheye2	150°	1280 × 1024	query
$c_3$	fisheye3	195°	1280 × 1024	query
$c_4$	pinhole	85°	1920 × 1200	query

Table 2. The representation and parameters of 5 cameras.

Scene	# Frames	# Frames Query (day / night)					Spatial Extent (m)
	Reference 360	360	Pinhole	Fisheye1	Fisheye2	Fisheye3	
Concourse	491	593/514	1186/1028	1186/1028	1186/1028	1186/1028	93 × 15
Hall	540	1123/1061	2246/2122	2246/2122	2246/2122	2246/2122	105 × 52
Atrium	581	875/1219	1750/2438	1750/2438	1750/2438	1750/2438	65 × 36
Patrium	632	1008/697	2016/1394	2016/1394	2016/1394	2016/1394	98 × 70

Table 3. 360Loc dataset description.

a FoV of  $360^\circ \times 30^\circ$ , angular resolution of  $0.2^\circ \times 2.0^\circ$ , and rotation rate of 10Hz, offering a comprehensive 360° environmental view. Regarding the calibration of the extrinsic poses between the lidar and the 360° camera, we employed a calibration toolbox [29] that applies to both lidar and camera projection models. This toolbox utilizes the SuperGlue [43] image matching pipeline to establish 2D-3D correspondences between the lidar and camera image. We perform pseudo-registration by synchronizing the two data modalities, images, and point clouds. Eventually, we use graph-based SLAM techniques for continuous pose estimations. In the four scenes, a total of 18 independent sequences of 360° images were captured (12 daytime, and 6 nighttime), resulting in a total number of 9334 images. For each scene, we selected a specific sequence captured during the daytime as the reference images, while the remaining images were defined as query images of the 360° image type. We provide more details and show why 360° mapping is superior to pinhole SfM in ambiguous scenes with repetitive and symmetric structures in the supplementary material.




Figure 3. Illustration of obtaining virtual camera images through random poses and image cropping.

### 3.1.1 Cross-device Queries

To enable a rigorous comparison of the difference in the performance of different FoV queries for visual localization tasks, we created four virtual cameras with diverse FoV from 360° cameras, which are shown in Figure 2. Given a 360° image  $\mathcal{I}_{c_0}$ , the corresponding virtual camera with pre-configured intrinsic parameters is extracted by

$$\mathcal{I}_{c_n} = \Psi_{c_n}(\mathcal{I}_{c_0}) = \pi_{c_n}^{-1}(\pi_{c_0}(\mathbf{R}\mathcal{I}_{c_0})), \quad (1)$$

where  $\pi_{c_n}$  denote the projection function of virtual camera and  $\pi_{c_0}$  is the projection function of 360° camera.  $\mathbf{R} \in SO(3)$  is a random relative rotation matrix to increase the diversity of views representing the scenes. Moreover, the inversed operation  $\Psi_{c_n}^{-1}$  can convert the  $c_n$  image back to a 360° image. As reported in Table 2, the virtual cameras include an undistorted pinhole model with 85° FoV and three fisheye cameras in Dual Sphere mode [56] with 120°, 150°, and 195° FoV respectively. Table 3 presents the number of image frames in the 360Loc dataset.

### 3.2. Ground Truth Generation

Besides the graph-based optimization in SLAM, we designed a set of offline optimization strategies to further improve the accuracy of camera pose estimation. After the acquisition of precise dense point cloud reconstructions and poses of 360° cameras, an Iterative Closest Point (ICP) algorithm is applied to align models between reference and the query sequences in the same scene. Moreover, we reconstructed the mesh model of the scenes and generated corresponding depth maps of 360° cameras.

**Bundle Adjustment (BA) of lidar mapping.** Incremental map construction can suffer from accumulating errors due to environmental degradation. We utilized a BA framework based on feature points extracted from lidar to refine the map and the poses. The optimization process involved minimizing the covariance matrix to constrain the distances between feature points and edge lines or plane features that are mutually matched.

First, we utilize an octree data structure to perform adaptive voxelization-based feature extraction. In this method, the point cloud map is segmented into voxels of predetermined size. Each voxel is checked to determine if its points  $P_u^f$  lie on a plane or a line, where  $u \in \{1, 2, \dots, U\}$ , obtained from the  $u$ -th frame of lidar scans. If not, the voxel is recursively subdivided using an octree structure until each voxel contains points  $P_u^f$  belonging to the same feature. Let's assume that the pose of the lidar in each frame is  $\eta = \{\eta_1, \eta_2, \dots, \eta_M\}$ , where  $\eta_u = (R_u, t_u | R_u \in SO(3), t_u \in \mathbb{R}^3)$ . In that case, the feature points in the global map can be represented as follows:

$$P_u = R_u \times P_u^f + t_u. \quad (2)$$

After simplifying the lidar map to edge or plane features, the process of BA becomes focused on determining the pose  $\eta$  and the location of the single feature, which can be represented as  $(n_f, q)$ , where  $q$  represents the location of a specific feature,  $n_f$  is the direction vector of an edge line or the normal vector of a plane. To minimize the distance between each feature point and the corresponding feature, we can utilize the BA:

$$(\eta^*, n_f^*, q^*) = \arg \min_{\eta, n_f, q} \frac{1}{U} \sum_{u=1}^U (n_f^T (P_u - q))^2. \quad (3)$$

It has been proved that when the plane's normal vector is set to the minimum eigenvector, and  $q$  is set to the centroid of the feature, i.e.  $q = \hat{P} = \frac{1}{U} \sum_{u=1}^U P_u$ , Eq. 3 reaches its minimum value. Additionally, the BA problem in lidar mapping has a closed-form solution that is independent of the features  $(n_f, q)$  [34]. It can be simplified to the following problem:

$$\eta^* = \arg \min_{\eta} \lambda_{\min} (A), \quad (4)$$

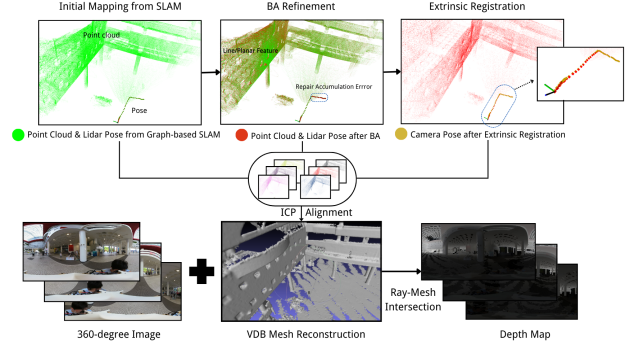


Figure 4. Overview of GT generation.

where,  $\lambda$  represents the eigenvalue of  $A$ , and

$$A = \frac{1}{U} \sum_{u=1}^U (P_u - \hat{P}) (P_u - \hat{P})^T. \quad (5)$$

Now, the BA problem is simplified by adjusting the lidar pose  $\eta$  to minimize the smallest eigenvalue  $\lambda_3$  of the point covariance matrix  $A$  defined in Eq. 5. By employing this strategy, we refined the pose  $\eta$  of each frame and the edge or plane features in the lidar map.

**Refined cameras poses.** The poses of 360° camera obtained from online SLAM are further optimized by the registration with respect to the dense refined point cloud model. Taking the pre-calibrated extrinsic parameters as the initial guess, we used the RANSAC to refine the lidar-camera transformation [29]. This registration process is based on the normalized information distance (NID) [51], which serves as a mutual information-based cross-modal distance metric. Finally, we align the reference models and query models into the same coordinate system to generate the ground truth for the query sequences. Specifically, we utilize the CloudCompare toolbox [19] to manually select feature points across multiple point cloud models as initial values. Then, we employ the ICP algorithm to register the point cloud models together. Afterwards, we employed a practical approach to volumetric surface reconstruction called Truncated Signed Distance Functions (TSDFs) [57] to achieve the reconstruction from point clouds to meshes with an efficient and sparse data structure called Voxel Data Base (VDB) [39]. At this stage, we can utilize the ray-mesh intersection method [15] to cast rays from cameras onto the mesh model. By intersecting the rays with the mesh, we can determine the depths of the corresponding points on the mesh surface. After a series of joint optimizations between multiple modalities, we have generated a set of GT data. Figure 2 shows some instances. This GT data includes reference images  $\mathcal{I}_{c_0}^r$ , the depth maps  $D_{c_0}^r$  of the reference images, and the reference maps containing the point cloud models  $\mathcal{P}$ , mesh models  $M$ , as well as camera pose odom-

entry  $\{\xi\}$ . Figure 4 summarizes the GT generation.

## 4. Omnidirectional Visual Localization

We extend the current feature-matching-based and absolute pose regression pipelines for omnidirectional visual localization. Given a query image  $\mathcal{I}^q$  in any camera model, we seek to estimate its poses within the environment modeled by  $360^\circ$  images  $\mathbf{I}^r$ . To minimize the domain gap between the query image from  $c_1, c_2, c_3, c_4$  and reference images, we explore visual cameras (VC) in two ways: VC1, remapping query images to  $360^\circ$  domain using  $\Psi_{c_n}^{-1}$ ; VC2, rectifying  $360^\circ$  images into queries’ domains using  $\Psi_{c_n}$ .

### 4.1. Feature-matching-based Localization

Most feature-matching-based techniques first perform IR to reduce the search space before estimating the pose.

#### 4.1.1 Image Retrieval

For method VC1, if query  $\mathcal{I}^q$  captured from  $c_0$ , we retrieve the  $k$  most similar images from  $\mathbf{I}^r$  by calculating and sorting  $\text{simi}_{\text{cos}}(\mathcal{F}(\mathcal{I}^q), \mathcal{F}(\mathcal{I}^r))$ ,  $\mathcal{I}^r \in \mathbf{I}^r$  and  $\mathcal{F}(\cdot)$  denotes the function to map each image to the global feature domain.  $\text{simi}_{\text{cos}}(\cdot)$  is cosine similarity for two feature embeddings. If query  $\mathcal{I}^q$  captured from  $c_1, c_2, c_3, c_4$ , we then retrieve top- $k$  reference images based on  $\text{simi}_{\text{cos}}(\mathcal{F}(\Psi_{c_n}^{-1}(\mathcal{I}^q)), \mathcal{F}(\mathcal{I}^r))$ ,  $\mathcal{I}^r \in \mathbf{I}^r$ .

In method VC2, we expand the global features for each  $360^\circ$  reference image by cameras  $c$  including virtual pinhole cameras forming a cube map and virtual fisheye cameras. We define the similarity score between  $\mathcal{I}^q$  and  $\mathcal{I}^r$  as:

$$\max(\text{simi}_{\text{cos}}(\mathcal{F}(\mathcal{I}^q), \mathcal{G}_{\mathcal{F}}(\mathcal{I}^r)), \quad (6)$$

where global feature group of reference is  $\mathcal{G}_{\mathcal{F}}(\mathcal{I}^r) = \{\mathcal{F}(\Psi_c(\mathcal{I}^r)), \dots\}$ . We use the highest similarity value calculated from  $\mathcal{F}(\mathcal{I}^q)$  and  $\mathcal{G}_{\mathcal{F}}(\mathcal{I}^r)$  as the similarity score for each  $\mathcal{I}^r$  to ensure retrieve  $k$  most similar  $360^\circ$  reference images because some rectified images are from the same  $\mathcal{I}^r$ . Note that we can eliminate the domain gap during the image retrieval step in this way.

#### 4.1.2 Local Feature Matching and Pose Estimation

For each pinhole query frame, we retrieve relevant reference images, match their local features, leverage the depth maps  $D_{c_0}$  to establish the 2D-3D correspondences, and finally estimate a pose with PnP+RANSAC. Unlike [11, 61], we directly match query image with retrieved  $360^\circ$  reference images described in Section 4.1.1. For query images from  $c_0, c_1, c_2, c_3$ , i.e., fisheye and  $360^\circ$  query frames, we utilize the function that calculates pose error in sphere camera model in OpenGV [28] library for PnP+RANSAC.

## 4.2. Absolute Pose Regression

APRs train deep neural networks to regress the 6DoF camera pose of a query image.

**PN.** PoseNet (PN) is the first APR model. Since there is no open source code [23, 24], we follow the modification in [8, 36] and use ResNet34 [21] as the backbone network. **MS-T.** MS-Transformer [49] is an APR model incorporating attention and implementing transformers as backbone. **We note APR methods using our virtual camera method, VC2, as APR<sup>vc2</sup>.** The difference between APR and APR<sup>vc2</sup> is the training stage. For APR baselines, the training set is  $\mathbf{I}^r$ . For APR<sup>vc2</sup>, they are trained with  $360^\circ$  images, cropped pinhole images, and cropped fisheye images, i.e.,  $\mathbf{I}^r \cup \Psi_c(\mathbf{I}^r)$  introduced in Section 4.1.1 and Eq. 1.

All APR models are implemented in Python using PyTorch [41]. During training, all input images are resized to  $256 \times 256$  and then randomly cropped to  $224 \times 224$ . For both PN and MS-T, we set an initial learning rate of  $\lambda = 10^{-4}$  and a batch size of 32 for 300 epochs of each scene. Training and evaluation in Section 5 are performed on an NVIDIA GeForce GTX 3090 GPU.

## 5. Evaluation

We provide detailed results for each scene in the dataset and more settings in supplementary material.

### 5.1. Image Retrieval

We evaluate global descriptors computed by NetVLAD [1], CosPlace [3], OpenIBL [18] and AP-GeM [20]. The query image is deemed correctly localized if at least one of the top  $k$  retrieved database images is within  $d = 5m$  from the ground truth position of the query for Concourse and  $d = 10m$  for the other three scenes. The image retrieval results are shown in Table 4. Among all global feature descriptor methods, the  $360^\circ$  query exhibits the best precision and recall in most cases, while the pinhole query performs the worst. The remap method (VC1) provides limited improvement for pinhole queries but yields higher improvement for fisheye1, fisheye2, and fisheye3 queries. The reason is that the FoV of pinhole cameras is only  $85^\circ$ . Consequently, VC1 results in significant black borders when converting to a  $360^\circ$  image due to the limited coverage.

The rectify method (VC2) significantly improves pinhole, fisheye1, fisheye2, and fisheye3 queries by eliminating the domain gap in IR. However, the pinhole, fisheye1, and fisheye2 queries’ recall and precision are still much lower than those of the  $360^\circ$  query. Only the query from fisheye3 (widest FoV) approaches the performance of  $360^\circ$  query. The domain gap mainly affects the precision and recall of fisheye3. Both remap (VC1) and crop (VC2) significantly improve IR performance for fisheye3. On the other hand, pinhole queries are more prone to being mistaken as erro-

Query	NetVLAD [1]					Cosplace [3]					OpenIBL [18]					AP-GeM [20]				
	R@1	R@5	P@5	R@10	P@10	R@1	R@5	P@5	R@10	P@10	R@1	R@5	P@5	R@10	P@10	R@1	R@5	P@5	R@10	P@10
pinhole	0.23	0.45	0.22	0.58	0.22	0.15	0.26	0.15	0.33	0.15	0.18	0.36	0.18	0.48	0.18	0.2	0.37	0.2	0.47	0.2
+VC1	0.24	0.45	0.24	0.57	0.23	0.21	0.33	0.21	0.41	0.21	0.21	0.39	0.21	0.5	0.2	0.25	0.42	0.25	0.53	0.24
+VC2	0.5	0.67	0.48	0.75	0.47	0.32	0.41	0.32	0.48	0.31	0.51	0.67	0.49	0.75	0.47	0.5	0.68	0.49	0.77	0.47
fisheye1	0.42	0.67	0.41	0.77	0.39	0.28	0.43	0.28	0.52	0.28	0.37	0.58	0.36	0.69	0.34	0.35	0.55	0.34	0.66	0.33
+VC1	0.51	0.72	0.49	0.8	0.47	0.36	0.48	0.35	0.56	0.34	0.52	0.7	0.5	0.79	0.48	0.43	0.62	0.42	0.72	0.4
+VC2	0.73	0.91	0.63	0.95	0.57	0.63	0.85	0.51	0.92	0.43	0.74	0.91	0.62	0.95	0.54	0.65	0.88	0.57	<b>0.94</b>	0.51
fisheye2	0.45	0.7	0.44	0.8	0.42	0.3	0.46	0.31	0.55	0.31	0.41	0.62	0.4	0.73	0.38	0.38	0.59	0.36	0.68	0.35
+VC1	0.54	0.74	0.52	0.83	0.49	0.37	0.49	0.36	0.57	0.35	0.56	0.73	0.54	0.81	0.51	0.46	0.65	0.45	0.74	0.43
+VC2	0.74	0.92	0.65	0.95	0.58	0.64	0.87	0.53	0.93	0.45	0.76	0.92	0.65	<b>0.96</b>	0.56	0.67	0.89	0.58	<b>0.94</b>	0.52
fisheye3	0.57	0.79	0.55	0.86	0.52	0.4	0.56	0.4	0.65	0.4	0.53	0.74	0.51	0.83	0.49	0.45	0.66	0.43	0.75	0.41
+VC1	0.63	0.81	0.61	0.88	0.58	0.48	0.61	0.48	0.68	0.47	0.67	0.82	0.65	0.88	0.61	0.55	0.73	0.53	0.81	0.51
+VC2	0.77	<b>0.93</b>	0.68	<b>0.96</b>	0.61	0.69	0.89	0.58	0.94	0.5	0.79	0.93	0.68	<b>0.96</b>	0.6	0.67	<b>0.9</b>	0.59	<b>0.94</b>	0.54
360	<b>0.79</b>	0.86	<b>0.77</b>	0.88	<b>0.73</b>	<b>0.92</b>	<b>0.95</b>	<b>0.91</b>	<b>0.96</b>	<b>0.89</b>	<b>0.89</b>	<b>0.94</b>	<b>0.88</b>	0.95	<b>0.83</b>	<b>0.79</b>	<b>0.9</b>	<b>0.77</b>	<b>0.94</b>	<b>0.72</b>

Table 4. Image retrieval results based on 360° reference database average over four scenes, the recall, and precision for the top  $k$  retrieved images,  $k = 1, 5, 10$ . # indicates the highest value of R@ $k$  and P@ $k$  for each device w and w/o virtual cameras (VC1, VC2). Best results for all devices of R@ $k$  and P@ $k$  are in bold with #.

	NetVLAD [1]						CosPlace [3]					
	DISK + LG		SP + LG		SP + SG		DISK + LG		SP + LG		SP + SG	
	Day	Night	Day	Night	Day	Night	Day	Night	Day	Night	Day	Night
pinhole	6.0/11.3/24.6	1.74/4/10.3	8.0/14.9/30.9	2.2/5.5/13.5	8.4/15.2/30.7	2.3/5.6/12.3	4.2/7.8/18.0	1.6/3.5/8.6	4.8/10.2/22.1	1.9/4.7/11.1	5.4/10.4/21.1	2.1/4.7/10.4
+VC1	8.5/14.0/23.5	2.24/1/7.9	10.4/17.0/27.5	2.9/5.3/10.1	10.9/17.8/28.5	2.8/5.6/9.9	6.1/10.8/21.1	1.7/3.6/8.2	7.5/13.2/22.5	2.0/4.5/9.6	7.6/13.5/22.8	2.14/7.9/6
+VC2	14.2/22.2/35.5	4.1/7.8/13.6	<b>19.8</b> / 29.7/42.9	6.1/10.4/16.9	<b>21.6/33.2</b> / 49.7	5.9 / 11.0 / 18.4	8.0/13.1/23.5	2.5/4.6/9.1	10.7/16.4/26.6	3.0/5.7/11.4	11.6/18.5/30.5	3.5/6.8/12.8
fisheye1	1.6/4.4/17.7	0.5/1.8/7.4	1.9/5.4/20.1	0.7/2.3/10.5	1.6/4.7/18.4	0.5/1.9/8.2	0.8/2.5/11.8	0.4/1.4/5.8	1.0/3.5/13.0	0.5/1.4/8.2	0.9/3.4/12.1	0.3/1.4/7.0
+VC1	3.3/9.2/27.6	0.8/2.7/9.6	4.1/10.6/32.2	1.4/4.4/14.9	3.0/9.5/29.6	0.9/3.1/11.7	2.3/5.5/19.4	0.5/1.6/7.3	2.1/6.1/19.9	0.7/2.2/9.0	1.9/5.5/19.1	0.5/1.9/7.3
+VC2	3.9/10.5/33.0	1.0/4.0/14.6	4.3/12.4/38.2	1.9/6.4/21.8	3.6/11.0/34.5	1.1/5.3/19.4	2.5/6.9/25.3	0.8/2.8/12.2	2.8/8.2/29.0	1.3/4.6/18.0	2.1/7.1/26.7	1.0/4.0/16.2
fisheye2	1.6/4.9/20.9	0.5/2.0/8.7	1.9/6.7/23.2	0.8/3.0/11.8	1.7/5.2/19.5	0.7/2.5/9.9	1.3/3.5/14.2	0.4/1.6/6.9	1.2/3.8/15.2	0.5/1.5/9.1	1.2/3.9/12.9	0.6/1.6/7.2
+VC1	4.3 /10.8/30.9	0.8/3.0/11.2	4.7/12.4/34.1	1.8/5.4/15.8	4.1 /10.6/31.5	1.1/3.6/13.7	2.5/6.5/20.6	0.5/1.7/7.4	2.5/7.0/22.1	0.8/2.4/9.4	2.2/6.8/20.2	0.5/2.1/8.0
+VC2	4.3/11.0/34.4	1.1/4.7/17.3	5.1/14.0/41.1	2.0/7.2/24.8	3.7 / 11.5/36.8	1.5/5.9/21.2	2.8/7.3/27.1	0.8/2.9/13.4	2.9/8.9/32.0	1.6/5.3/20.1	2.5/8.0/27.9	1.1/4.2/17.7
fisheye3	3.8/9.5/29.8	1.0/3.6/13.8	4.0/10.5/31.6	1.3/4.6/16.4	3.4/9.1/28.4	0.8/3.8/13.8	2.5/6.3/21.9	0.6/2.4/10.1	2.8/7.2/22.3	0.9/2.9/12.4	2.0/5.9/20.0	1.3/4.2/15.0
+VC1	5.9/14.7 /39.5	1.5/5.2/17.7	6.0 /16.2/43.5	2.0/6.8/21.9	5.8/14.7 /39.1	1.8/5.5/18.3	4.4/10.2 /30.1	1.1/3.3/12.8	4.6/11.6/32.0	1.4/4.1/14.4	4.3/10.5 /29.7	1.2/3.8/12.3
+VC2	5.2/13.9 /41.8	2.1/6.5/22.5	5.9 /16.5/46.3	2.5/8.6/29.1	5.4/14.2 /40.5	2.1/7.3/25.9	4.3/9.8 /34.6	1.7/5.2/19.5	4.7/12.6 /36.8	2.2/7.1/23.8	3.8 / 10.5 /32.5	1.6/5.1/20.7
360	<b>17.1 / 30.8 / 66.1</b>	<b>8.5 / 20.1 / 47.5</b>	18.2 / 34.6 / 64.2	<b>7.0 / 18.7 / 45.3</b>	15.8 / 31.2 / 60.4	<b>7.0 / 17.8 / 42.8</b>	<b>17.6 / 31.8 / 68.1</b>	<b>8.7 / 22.0 / 56.0</b>	<b>18.7 / 34.9 / 68.1</b>	<b>7.3 / 20.0 / 53.4</b>	<b>16.6 / 32.6 / 65.7</b>	<b>7.1 / 18.7 / 50.4</b>

Table 5. Local matching localization results. The average percentage of predictions with high (0.25m, 2°), medium (0.5m, 5°), and low (5m, 10°) accuracy [46] (higher is better) over four scenes. # indicates the highest value for each device w and w/o virtual cameras (VC1, VC2) of each accuracy level. The best results for all devices of each accuracy level are in bold with #.

neous locations with similar structures due to their narrower FoV even there is no cross-device domain gap during IR by applying VC2 (Some figures in supplementary material).

## 5.2. Visual Localization

We compare our approach with the following baselines in two categories: 1) Local feature matching pipelines tailored from HLoc [42], using different keypoint descriptors (Superpoint (SP) [16] and DISK [55]), and matchers (Super-Glue (SG) [43], follow-up SOTA LightGlue (LG) [31]). 2) The end-to-end APRs: PN [23, 24] and MS-T [49].

**Local feature matching:** During local feature matching, all 360° images are cropped to 1228 × 614 because of the tradeoff of time and computation. We report the average results over four scenes in Table 5. The 360° query achieves the best performance in three accuracy levels in most cases

across all IR, keypoint descriptors, and matchers settings. It is especially more robust in challenging nighttime conditions. VC1 and VC2 techniques improve the recall and precision of IR, increasing the accuracy of 2D-2D matching for all cameras. In most cases, the performance at the low accuracy level (5m, 10°) is correlated with the FoV, where a larger FoV results in higher performance. However, the pinhole query with VC2 during IR performs comparably to the 360° queries at the high (0.25m, 2°) and median (0.5m, 5°) accuracy levels. In contrast, query frames from  $c_1$ ,  $c_2$  and  $c_3$  demonstrate relatively lower performance at the high and medium accuracy levels.

As observed in Table 4, different IR methods display different performances depending on the type of camera. We thus consider both NetVLAD and CosPlace in visual localization. In most cases, 360° query frames achieve higher



accuracy with CosPlace while pinhole and fisheye query frames have lower accuracy than NetVLAD as shown in Table 5. These results match the precision and recall difference noted in Table 4. We believe that the FoV not only affects the robustness of IR but also has an impact on local 2D-2D matching performance. Pinhole queries suffer from erroneous matches due to interference from symmetrical and repetitive structures, while the larger FoV of fisheye and 360° query frames capture more unique visual features. We provide examples in the supplementary material.

**APR:** APRs cannot extrapolate well beyond the training set [40, 47]. cross-device queries further complicate this challenge by introducing an additional dimension of FoV. Due to the high efficiency of 360° mapping, the training set  $\mathbf{I}^r$  in 360Loc contains only around one-third of the images compared to datasets [24]. Figure 5 shows that when PN and MS-T are trained solely on  $\mathbf{I}^r$  with only 360° images, a smaller domain gap between the query and the 360° image yields a lower error. However, when we introduce images from virtual cameras for data augmentation,  $\text{PN}^{vc2}$  and  $\text{MS-T}^{vc2}$  exhibit significantly reduced translation and rotation errors across all queries, particularly during daytime.  $\text{MS-T}^{vc2}$  reduces translation error by up to 79% and rotation error by up to 72% compared to MS-T.  $\text{PN}^{vc2}$  displays similar improvement over PN. In most cases, except for  $\text{PN}^{vc2}$ 's rotation error for the 360° queries during daytime, both the 360° and fisheye queries exhibit higher accuracy than the pinhole query on  $\text{PN}^{vc2}$  and  $\text{MS-T}^{vc2}$ . This suggests that a larger FoV still helps improve visual localization accuracy in challenging scenes. Another interesting finding is that even though the augmented training set  $\mathbf{I}^r \cup \Psi_c(\mathbf{I}^r)$ , which includes virtual camera images, does not increase the number of 360° images, the error for the 360° query still decreases. This reduction is particularly noticeable in the case of translation errors during daytime. The result fully demonstrates the utility of employing virtual cameras for data augmentation.

### 5.3. Analysis

Cross-device visual positioning presents significant challenges for IR, local matching, and APRs. Our VC1 and VC2 methods demonstrate practical enhancements in the performance of IR and APR for cross-device scenarios. However, it is essential to note that during the local matching process, the accuracy of matches and the recall and precision of IR for query frames from different cameras may not align perfectly. The chosen IR method and its training noticeably affect accuracy for similar cameras. Fisheye cameras exhibit better performance in IR compared to pinhole cameras. However, pinhole cameras outperform fisheye cameras for high accuracy and median accuracy levels in local matching. This is likely due to existing feature extraction and matching models lacking training data on 360° and fisheye

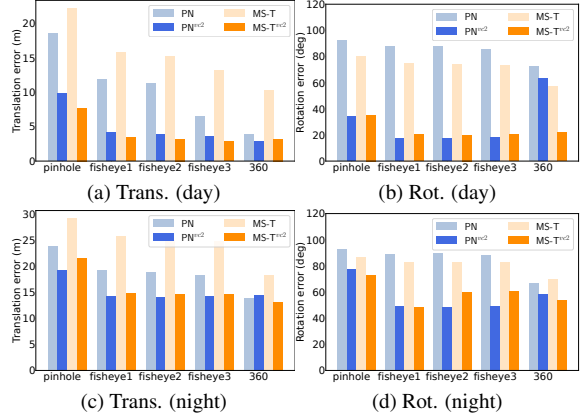


Figure 5. The average of median translation/rotation errors in ( $m/^\circ$ ) over 4 scenes.

cameras, resulting in less accurate matching. We attribute the inferior performance of pinhole query frames at the low accuracy level to IR’s insufficient recall and precision. Additionally, pinhole queries are more susceptible to interference when there are many repetitive and symmetrical features in the scene, even when the retrieved reference image is correct (some example figures in the supplementary material). By utilizing VC2 to augment IR and APR’s training data, we eliminate the cross-device domain gap. We demonstrate that panoramic perspective and a larger FoV can significantly improve the performance of IR and APRs and find that query frames from 360° camera and ultra-wide FoV cameras are less prone to being misidentified as erroneous locations with similar structures. This result suggests the promising potential of fisheye and 360° cameras as viable sensors for localization tasks in indoor environments with low GPS accuracy.

## 6. Conclusion

360Loc is the first dataset and benchmark that explores the challenge of cross-device visual positioning, involving 360° reference frames, and query frames from pinhole, ultra-wide FoV fisheye, and 360° cameras. We first identified the absence of datasets with ground truth 6DoF poses for 360° images, and the limited research on cross-device localization and the robustness of different cameras in ambiguous scenes. To address these limitations, we build a dataset with 360° images as reference and query frames from pinhole, ultra-wide FoV fisheye camera and 360° cameras via a *virtual camera* solution. This method enables fair comparisons in cross-device visual localization tasks and helps reduce the domain gap between different cameras. By evaluating feature-matching-based and pose regression-based methods, we demonstrate the effectiveness of our virtual camera approach and the increased robustness of 360° cameras in visual localization for challenging and ambiguous scenes.



## References

- [1] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Netvlad: Cnn architecture for weakly supervised place recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5297–5307, 2016. 2, 6, 7
- [2] Iro Armeni, Sasha Sax, Amir R Zamir, and Silvio Savarese. Joint 2d-3d-semantic data for indoor scene understanding. *arXiv preprint arXiv:1702.01105*, 2017. 3
- [3] Gabriele Berton, Carlo Masone, and Barbara Caputo. Rethinking visual geo-localization for large-scale applications. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4878–4888, 2022. 2, 6, 7
- [4] Hunter Blanton, Connor Greenwell, Scott Workman, and Nathan Jacobs. Extending absolute pose regression to multiple scenes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020. 3
- [5] Eric Brachmann and Carsten Rother. Learning less is more - 6D camera localization via 3D surface regression. In *CVPR*, 2018. 1, 2
- [6] Eric Brachmann and Carsten Rother. Visual camera relocalization from RGB and RGB-D images using DSAC. *TPAMI*, 2021.
- [7] Eric Brachmann, Alexander Krull, Sebastian Nowozin, Jamie Shotton, Frank Michel, Stefan Gumhold, and Carsten Rother. DSAC-Differentiable RANSAC for camera localization. In *CVPR*, 2017. 1, 2
- [8] Samarth Brahmabhatt, Jinwei Gu, Kihwan Kim, James Hays, and Jan Kautz. Geometry-aware learning of maps for camera localization. In *IEEE conference on computer vision and pattern recognition*, 2018. 3, 6
- [9] Mai Bui, Tolga Birdal, Haowen Deng, Shadi Albarqouni, Leonidas Guibas, Slobodan Ilic, and Nassir Navab. 6d camera relocalization in ambiguous scenes via continuous multi-modal inference. 2020. 3, 4
- [10] Nicholas Carlevaris-Bianco, Arash K Ushani, and Ryan M Eustice. University of michigan north campus long-term vision and lidar dataset. *The International Journal of Robotics Research*, 35(9):1023–1035, 2016. 3, 4
- [11] David M Chen, Georges Baatz, Kevin Köser, Sam S Tsai, Ramakrishna Vedantham, Timo Pylvänäinen, Kimmo Roimela, Xin Chen, Jeff Bach, Marc Pollefeys, et al. City-scale landmark identification on mobile devices. In *CVPR 2011*, pages 737–744. IEEE, 2011. 3, 4, 6
- [12] Shuai Chen, Zirui Wang, and Victor Prisacariu. Direct-posenet: absolute pose regression with photometric consistency. In *2021 International Conference on 3D Vision (3DV)*, pages 1175–1185. IEEE, 2021. 3
- [13] Shuai Chen, Xinghui Li, Zirui Wang, and Victor A Prisacariu. Dfnet: Enhance absolute pose regression with direct feature matching. In *ECCV 2022. Tel Aviv, Israel, October 23–27, 2022, Part X*. Springer, 2022. 3
- [14] Santiago Cortés, Arno Solin, Esa Rahtu, and Juho Kannala. Advio: An authentic dataset for visual-inertial odometry. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 419–434, 2018. 3, 4
- [15] Dawson-Haggerty et al. trimesh. 5
- [16] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 224–236, 2018. 2, 7
- [17] Mihai Dusmanu, Ignacio Rocco, Tomas Pajdla, Marc Pollefeys, Josef Sivic, Akihiko Torii, and Torsten Sattler. 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*, pages 8092–8101, 2019. 1
- [18] Yixiao Ge, Haibo Wang, Feng Zhu, Rui Zhao, and Hongsheng Li. Self-supervising fine-grained region similarities for large-scale image localization. In *European Conference on Computer Vision*, 2020. 2, 6, 7
- [19] Daniel Girardeau-Montaut. Cloudcompare. *France: EDF R&D Telecom ParisTech*, 11, 2016. 5
- [20] A. Gordo, J. Almazan, J. Revaud, and D. Larlus. End-to-end learning of deep visual representations for image retrieval. *IJCV*, 2017. 2, 6, 7
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [22] Huajian Huang, Yinzhe Xu, Yingshu Chen, and Sai-Kit Yeung. 360vot: A new benchmark dataset for omnidirectional visual object tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20566–20576, 2023. 1
- [23] Alex Kendall and Roberto Cipolla. Geometric loss functions for camera pose regression with deep learning. In *IEEE conference on computer vision and pattern recognition*, pages 5974–5983, 2017. 1, 3, 6, 7
- [24] Alex Kendall, Matthew Grimes, and Roberto Cipolla. Posenet: A convolutional network for real-time 6-dof camera relocalization. In *Proceedings of the IEEE international conference on computer vision*, pages 2938–2946, 2015. 1, 3, 4, 6, 7, 8
- [25] Junho Kim, Changwoon Choi, Hojun Jang, and Young Min Kim. Piccolo: point cloud-centric omnidirectional localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3313–3323, 2021. 2, 3
- [26] Junho Kim, Hojun Jang, Changwoon Choi, and Young Min Kim. Cpo: Change robust panorama to point cloud localization. In *European Conference on Computer Vision*, pages 176–192. Springer, 2022.
- [27] Junho Kim, Eun Sun Lee, and Young Min Kim. Calibrating panoramic depth estimation for practical localization and mapping. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8830–8840, 2023. 2
- [28] Laurent Kneip and Paul Furgale. Opengv: A unified and generalized approach to real-time calibrated geometric vision. In *2014 IEEE international conference on robotics and automation (ICRA)*, pages 1–8. IEEE, 2014. 6

- [29] Kenji Koide, Shuji Oishi, Masashi Yokozuka, and Atsuhiko Banno. General, single-shot, target-less, and automatic lidar-camera extrinsic calibration toolbox. *arXiv preprint arXiv:2302.05094*, 2023. 4, 5
- [30] Donghwan Lee, Soohyun Ryu, Suyong Yeon, Yonghan Lee, Deokhwa Kim, Cheolho Han, Yohann Cabon, Philippe Weinzaepfel, Nicolas Guérin, Gabriela Csurka, et al. Large-scale localization datasets in crowded indoor spaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3227–3236, 2021. 3, 4
- [31] Philipp Lindenberger, Paul-Edouard Sarlin, and Marc Pollefeys. LightGlue: Local Feature Matching at Light Speed. In *ICCV*, 2023. 7
- [32] Haomin Liu, Linsheng Zhao, Zhen Peng, Weijian Xie, Mingxuan Jiang, Hongbin Zha, Hujun Bao, and Guofeng Zhang. A low-cost and scalable framework to build large-scale localization benchmark for augmented reality. *IEEE Transactions on Circuits and Systems for Video Technology*, 2023. 3
- [33] Liu Liu, Hongdong Li, and Yuchao Dai. Efficient global 2d-3d matching for camera localization in a large-scale 3d map. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2372–2381, 2017. 1
- [34] Zheng Liu and Fu Zhang. Balm: Bundle adjustment for lidar mapping. *IEEE Robotics and Automation Letters*, 6(2): 3184–3191, 2021. 5
- [35] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60:91–110, 2004. 2
- [36] Iaroslav Melekhov, Juha Ylioinas, Juho Kannala, and Esa Rahtu. Image-based localization using hourglass networks. In *IEEE international conference on computer vision workshops*, 2017. 3, 6
- [37] Arthur Moreau, Nathan Piasco, Dzmityr Tsishkou, Bogdan Stanciulescu, and Arnaud de La Fortelle. Coordinet: uncertainty-aware pose regressor for reliable vehicle localization. In *IEEE/CVF Winter Conference on Applications of Computer Vision*, 2022. 3
- [38] Jeffri Murrugarra-Llerena, Thiago LT Da Silveira, and Claudio R Jung. Pose estimation for two-view panoramas based on keypoint matching: A comparative study and critical analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5202–5211, 2022. 3
- [39] Ken Museth. Vdb: High-resolution sparse volumes with dynamic topology. *ACM transactions on graphics (TOG)*, 32(3):1–22, 2013. 5
- [40] Tony Ng, Adrian Lopez-Rodriguez, Vassileios Balntas, and Krystian Mikolajczyk. Reassessing the limitations of cnn methods for camera pose regression. *arXiv preprint arXiv:2108.07260*, 2021. 8
- [41] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 6
- [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*, pages 12716–12725, 2019. 1, 2, 7
- [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*, pages 4938–4947, 2020. 2, 4, 7
- [44] Paul-Edouard Sarlin, Mihai Dusmanu, Johannes L Schönberger, Pablo Speciale, Lukas Gruber, Viktor Larsson, Ondrej Miksik, and Marc Pollefeys. Lamar: Benchmarking localization and mapping for augmented reality. In *European Conference on Computer Vision*, pages 686–704. Springer, 2022. 3, 4
- [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*, 39(9):1744–1756, 2016. 1, 2, 3
- [46] Torsten Sattler, Will Maddern, Carl Toft, Akihiko Torii, Lars Hammarstrand, Erik Stenborg, Daniel Safari, Masatoshi Okutomi, Marc Pollefeys, Josef Sivic, et al. Benchmarking 6dof outdoor visual localization in changing conditions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8601–8610, 2018. 3, 4, 7
- [47] Torsten Sattler, Qunjie Zhou, Marc Pollefeys, and Laura Leal-Taixe. Understanding the limitations of cnn-based absolute camera pose regression. In *IEEE/CVF conference on computer vision and pattern recognition*, 2019. 3, 8
- [48] Thomas Schops, Johannes L Schonberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3260–3269, 2017. 3, 4
- [49] Yoli Shavit, Ron Ferens, and Yosi Keller. Learning multi-scene absolute pose regression with transformers. In *IEEE/CVF International Conference on Computer Vision*, pages 2733–2742, 2021. 1, 3, 6, 7
- [50] 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 conference on computer vision and pattern recognition*, pages 2930–2937, 2013. 3, 4
- [51] Alexander D Stewart. *Localisation using the appearance of prior structure*. PhD thesis, University of Oxford, 2014. 5
- [52] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. Loftr: Detector-free local feature matching with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8922–8931, 2021. 2
- [53] Xun Sun, Yuanfan Xie, Pei Luo, and Liang Wang. A dataset for benchmarking image-based localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7436–7444, 2017. 3, 4
- [54] Hajime Taira, Masatoshi Okutomi, Torsten Sattler, Mircea Cimpoi, Marc Pollefeys, Josef Sivic, Tomas Pajdla, and Ak-

- ihiko Torii. Inloc: Indoor visual localization with dense matching and view synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7199–7209, 2018. 1, 3, 4
- [55] Michał Tyszkiewicz, Pascal Fua, and Eduard Trulls. Disk: Learning local features with policy gradient. *Advances in Neural Information Processing Systems*, 33:14254–14265, 2020. 2, 7
- [56] Vladyslav Usenko, Nikolaus Demmel, and Daniel Cremers. The double sphere camera model. In *2018 International Conference on 3D Vision (3DV)*, pages 552–560. IEEE, 2018. 4
- [57] Ignacio Vizzo, Tiziano Guadagnino, Jens Behley, and Cyrill Stachniss. Vdbfusion: Flexible and efficient tsdf integration of range sensor data. *Sensors*, 22(3):1296, 2022. 5
- [58] Johanna Wald, Torsten Sattler, Stuart Golodetz, Tommaso Cavallari, and Federico Tombari. Beyond controlled environments: 3d camera re-localization in changing indoor scenes. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VII 16*, pages 467–487. Springer, 2020. 3, 4
- [59] Jian Wu, Liwei Ma, and Xiaolin Hu. Delving deeper into convolutional neural networks for camera relocalization. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5644–5651. IEEE, 2017. 3
- [60] Hang Xu, Qiang Zhao, Yike Ma, Xiaodong Li, Peng Yuan, Bailan Feng, Chenggang Yan, and Feng Dai. Pandora: A panoramic detection dataset for object with orientation. In *ECCV*, 2022. 1
- [61] Shen Yan, Yu Liu, Long Wang, Zehong Shen, Zhen Peng, Haomin Liu, Maojun Zhang, Guofeng Zhang, and Xiaowei Zhou. Long-term visual localization with mobile sensors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17245–17255, 2023. 1, 3, 4, 6
- [62] Dawen Yu and Shunping Ji. Grid based spherical cnn for object detection from panoramic images. *Sensors*, 19(11): 2622, 2019. 1