

Scene-agnostic Pose Regression for Visual Localization

Junwei Zheng¹ Ruiping Liu¹ Yufan Chen¹ Zhenfang Chen⁴ Kailun Yang³
 Jiaming Zhang^{1,2,*} Rainer Stiefelhagen¹

¹Karlsruhe Institute of Technology ²ETH Zurich ³Hunan University ⁴MIT-IBM Watson AI Lab

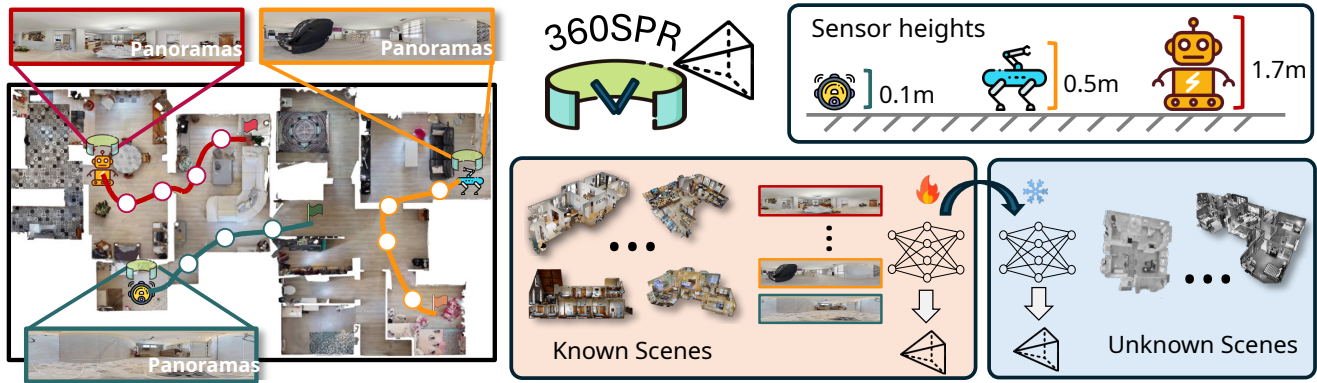


Figure 1. **360SPR** panoramic dataset for Scene-agnostic Pose Regression (SPR). Trajectories with different lengths and sampling intervals are collected at 3 heights of 360° sensors, *i.e.*, 0.1m, 0.5m, 1.7m, corresponding to sweeping 🧹, quadruped 🦘, and humanoid 🤖 robots. The proposed SPR model trained from known scenes can generalize well to unknown scenes, without database retrieval.

Abstract

*Absolute Pose Regression (APR) predicts 6D camera poses but lacks the adaptability to unknown environments without retraining, while Relative Pose Regression (RPR) generalizes better yet requires a large image retrieval database. Visual Odometry (VO) generalizes well in unseen environments but suffers from accumulated error in open trajectories. To address this dilemma, we introduce a new task, **Scene-agnostic Pose Regression (SPR)**, which can achieve accurate pose regression in a flexible way while eliminating the need for retraining or databases. To benchmark SPR, we created a large-scale dataset, **360SPR**, with over 200K photorealistic panoramas, 3.6M pinhole images and camera poses in 270 scenes at three different sensor heights. Furthermore, a **SPR-Mamba** model is initially proposed to address SPR in a dual-branch manner. Extensive experiments and studies demonstrate the effectiveness of our SPR paradigm, dataset, and model. In the unknown scenes of both 360SPR and 360Loc datasets, our method consistently outperforms APR, RPR and VO. The dataset and code are available at [SPR](#).*

*Corresponding author (e-mail: jiaming.zhang@kit.edu).

1. Introduction

Visual localization is one of the fundamental tasks in the computer vision domain. It is required to predict the 6D camera poses given the visual cues captured by a camera. The ability to determine the camera poses from captured images enables various downstream applications that require precise spatial awareness, *e.g.*, VR/AR [26, 54], autonomous driving [43], and robotics [48, 51].

Absolute Pose Regression (APR) is one of the classical paradigms for camera pose regression. Given the images of a scene, the model predicts the absolute camera poses concerning the scene coordinate system. Since a model in the APR paradigm only learns the scene-specific features, it is not applicable in unknown environments without retraining. Relative Pose Regression (RPR), in contrast, generalizes better in unknown environments. This paradigm learns the relative pose between a reference and query image during training. In the inference phase, the model first retrieves a reference image from the training set and then predicts the relative pose for a query image. Although RPR is able to predict camera poses in unseen environments, it requires a large database to retrieve reference images similar to the query image, which plays an important role in the whole RPR paradigm. Without enough overlap or similarity be-

tween reference and query images, the model performance of RPR drops dramatically. Visual Odometry (VO) predicts the next camera pose based on the previously predicted pose and can generalize well in unseen environments. However, it suffers from accumulated drift in open trajectories.

To address the dilemma of APR, RPR and VO, in this work, we introduce a novel task termed **Scene-agnostic Pose Regression (SPR)**, targeting the generalization problem and the need for a large database. Given a sequence of images along a trajectory, SPR takes the first image as the coordinate system origin. Using all images before the query image along the path as the model input, the camera pose of the query image is predicted *w.r.t.* the origin image, not the previous one. It's worth noting that there is no accumulated drift for open trajectories in the SPR paradigm compared to VO since the prediction of the current frame camera pose doesn't depend on the previous frame camera pose. Compared with APR, SPR disentangles the coordinate system from scenes and learns the relative poses between frames instead of scene-specific features so it is applicable in unseen environments. Unlike RPR, SPR bypasses the need to use a large database for image retrieval. The relative poses are calculated between the first frame of the trajectory and the desired query frames.

Most publicly available visual localization datasets utilize pinhole images to perform camera pose regression while only a small amount of datasets use panoramas, *e.g.*, the 360Loc [25] dataset. Fig. 2 illustrates the average median translation and rotation error of APR (represented by PoseNet [30]), RPR (represented by Relpose-GNN [63]) and SPR (represented by SPR-Mamba) with the change of image Field of View (FoV) in known environments of the 360Loc dataset. Both translation and rotation errors decrease when the image FoV increases, proving the necessity of utilizing panoramas in pose regression tasks. The reason behind this phenomenon is straightforward. Panoramas [24, 68, 79] not only enrich necessary visual information for camera pose regression but also enable sufficient overlap and similarity between frames, which is extremely important for RPR and SPR. The limitation of the panoramic visual localization dataset, 360Loc [25], is also obvious because it only provides less than 10K panoramas distributed in only 4 scenes with a fixed sampling interval and sensor height. The limited data capacity and diversity cannot meet the demand for accurate and robust spatial awareness in various real-world applications. As shown in Table 5a, the model trained on 360Loc [25] is not able to perform well in other diverse environments with an error increase of $4.08m/4.67^\circ$ on translation and rotation, respectively. Table 5b further demonstrates the lack of robustness when the model is trained at a fixed sensor height. To this end, we establish a large-scale dataset named **360SPR** with over 200K photorealistic panoramas and 3.6M pinhole im-




Figure 2. Model performance in APR, RPR, and SPR paradigms with the change of image field of view. TE and RE stand for Translation Error and Rotation Error, respectively.

ages with camera poses in 270 scenes using the Habitat simulator [47, 55, 59] powered by HM3D [49] and Matterport3D [6] datasets. Fig. 1 illustrates the data collection process. Trajectories of different lengths, with varying sampling intervals between sampling points along the path, are collected at three different robot heights, *i.e.*, 1.7m, 0.5m, 0.1m, corresponding to humanoid robots 🧑, quadruped robots 🐾, and sweeping robots 🧹, respectively.

Apart from the SPR task and 360SPR dataset, we propose a new model termed **SPR-Mamba** to explore the effectiveness of the SPR paradigm. SPR-Mamba consists of a local branch and a global branch. The local branch learns the relative pose between the current and previous frame while the global branch focuses on the one between the query and origin frame. These complementary branches enable comprehensive learning for camera pose regression. Extensive experiments on the 360SPR and 360Loc [25] benchmarks verify the effectiveness of the SPR with an error decrease over **7m/16°** on translation and rotation in unseen environments, outperforming APR and RPR paradigms in the camera pose regression task.

To summarize, our contributions are as follows:

- We propose a new task termed **Scene-agnostic Pose Regression (SPR)**, addressing the generalization problem of APR, the demand for a large database of RPR and the accumulated error of VO.
- We create **360SPR**, a dataset with 200K panoramas and 3.6M pinhole images across 270 scenes at three robot heights for panoramic visual localization.
- We introduce a model termed **SPR-Mamba** consisting of a local and global branch. The local branch learns the relative pose between the current and the previous image while the global branch focuses on the one between the query and the origin image.
- Extensive experiments on 360SPR and 360Loc [25] datasets prove the effectiveness of the proposed SPR paradigm and SPR-Mamba model with an error reduction over **7m/16°** on translation and rotation in unseen environments compared to APR and RPR.

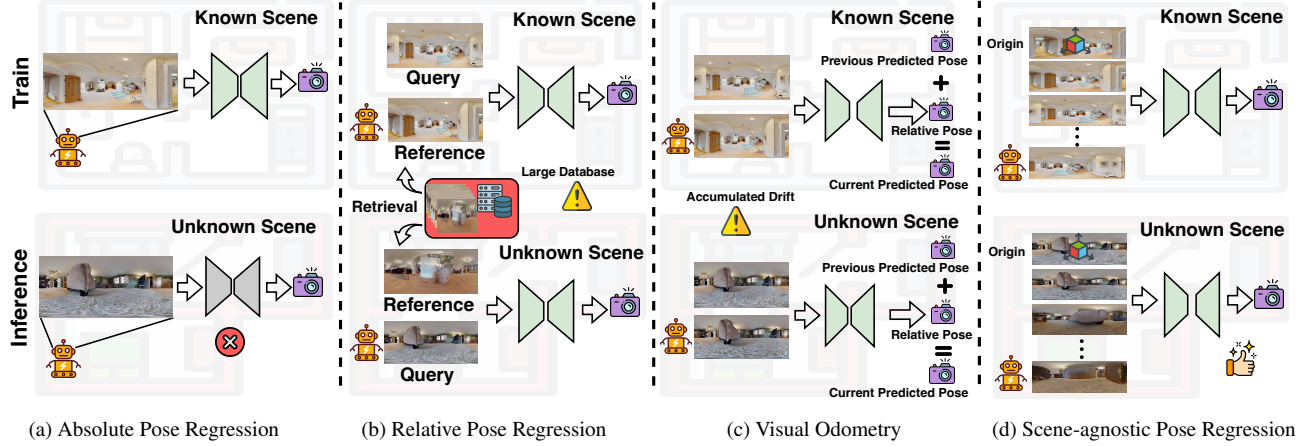


Figure 3. Comparison of different paradigms. (a) APR is not applicable in unseen scenes without retraining. (b) RPR requires a large database for image retrieval. (c) VO suffers from accumulated drift. (d) SPR (ours) sets the first image in a sequence as the origin and predicts the relative pose of its followings, *i.e.*, query images, avoiding the need of retraining, retrieval database or accumulated drift.

2. Related Work

Absolute Pose Regression (APR) methods [3, 8, 11, 28–30, 39, 52, 56, 69–71] aim to directly predict the 6D camera poses from images within a known environment used for training. PoseNet [30] employed a CNN to regress poses from single images. MS-Transformer [56] utilized transformer [65] for multiple scenes. AnchorPoint [52] introduced anchor points to enhance localization, while Marepo [8] conditioned the pose regressor on a scene-specific map representation and PMNet [39] leveraged neural volumetric pose features for localization. Besides, sequence-based APR methods leverage temporal dependency information from image sequences to improve absolute pose prediction performance. VidLoc [11] extended PoseNet by incorporating long short-term memory networks to capture temporal dependencies. MapNet [3], LS-G [70], and GL-Net [71] integrated geometric consistency and global-local feature fusion. While APR requires retraining on new scenes due to their scene-specific features, our SPR predicts camera poses in unknown scenes without retraining or dependence on scene-specific maps.






Relative Pose Regression (RPR) methods [2, 15, 22, 35, 36, 53, 63, 72, 74, 77, 80] estimate the 6D camera relative pose between the query image and reference image, generalizing to unseen scenes compared to APR methods. NN-Net [35] retrieved similar database images and predicted relative poses without requiring scene-specific training. CamNet [15] presented a coarse-to-fine retrieval framework that enhances retrieval accuracy. RelocNet [2] learned feature embeddings through continuous metric learning. EssNet [80] directly predicted the essential matrix from image pairs, facilitating robust RPR. RelPose-GNN [63] utilized a GNN to model spatial relationships between cameras. Although RPR methods demonstrate better generalization than APR methods in unseen environments, they



















normally require a large database for image retrieval and feature matching, which can be impractical for large-scale or dynamic settings. Additionally, the pose regression results may suffer from limited accuracy due to challenges in matching and environmental variations. Unlike RPR, SPR eliminates the reliance on large-scale reference image databases, yielding better pose prediction in unseen cases.

Visual Odometry (VO) is a core technique [12, 37, 38, 42, 45, 46, 57, 58, 61, 66, 76, 81] used in Simultaneous Localization and Mapping (SLAM) [4, 44] to estimate the trajectory of a camera by analyzing sequential image data to facilitate navigation in an unknown environment. Traditional feature-based VO methods [1, 16, 18, 31, 32] rely on identifying and tracking specific image features to estimate camera movement. Chien *et al.* [10] evaluated the choice of different features for monocular visual odometry. Learning-based VO methods leverage deep learning for visual localization. DeepVO [67] utilized neural networks to predict trajectories from image sequences, while D3VO [73] further added depth data for better accuracy. VRVO [78] was tailored for the smooth tracking demands of virtual environments. DAVO [33] was designed to quickly adapt to dynamic settings. CEGVO [27] employed a novel loss function to enhance visual odometry. ColVO [40] was proposed to estimate colon depth and colonoscopic pose continuously. The open-loop trajectory of VO has unavoidable drift accumulated over time and distance. In contrast, SPR is more robust in open- and close-loop trajectories.

3. Methodology

The scene-agnostic Pose Regression (SPR) task is formulated in Sec. 3.1 and compared with APR, RPR and VO. Besides, our large-scale 360SPR dataset for benchmarking the SPR task is presented in Sec. 3.2. In Sec. 3.3, we introduce the proposed dual-branch SPR-Mamba model.

Table 1. Comparison of visual localization datasets. Differences include: photographic  vs. photorealistic  types; field of views (panoramic vs. pinhole); image sizes; the number of scenes/trajectories/intervals; and sensor heights (in sweeping , quadruped , humanoid  robots). Our 360SPR has 200K panoramas and 3.6M pinhole images, offering diverse data for visual localization.

Dataset	Type	Panoramas	Pan. Res.	Pinholes	Pin. Res.	Scenes	Trajectories	Varying Interval	Height	Depth
7Scenes [19]		✗	✗	33K	480 × 640	7	?	?		✓
12Scenes [64]		✗	✗	246.7K	968 × 1296	12	?	?		✓
InLoc [60]		✗	✗	14K	1200 × 1600	5	?	?		✗
Cambridge [30]		✗	✗	13.2K	1080 × 1920	5	?	✗		✓
NaVIP [75]		✗	✗	300K	2160 × 3840	4	?	?		✗
LaMAR [54]		✗	✗	152K	480 × 640	3	?	?		✓
360Loc [25]		9.3K	3072 × 6144	14.2K	1200 × 1920	4	18	✗		✓
360SPR (ours)		200K	1024 × 2048	3.6M	512 × 512	270	20K	✓	  	✓




3.1. Scene-agnostic Pose Regression

Task Definition. The Scene-agnostic Pose Regression task aims to estimate the camera pose \mathbf{T}_q of a query image I_q relative to an origin image I_1 within an arbitrary scene, independent of specific scene characteristics or databases. Given a sequence of images I_1, I_2, \dots, I_q captured along a trajectory, the SPR task requires the model to accurately compute \mathbf{T}_q , which represents the camera’s position and orientation for image I_q . Unlike traditional pose estimation tasks that rely on structured scenes or predefined databases, SPR operates across diverse environments, using only the image sequence itself without scene-specific priors. The goal of SPR is to develop a robust model that outputs \mathbf{T}_q w.r.t. I_1 by processing the entire sequence I_1, \dots, I_{q-1}, I_q .

Task Difference. Fig. 3 showcases the difference between APR, RPR, VO and our proposed SPR task. ❶ APR (in Fig. 3a) is capable of predicting camera poses in seen environments occurring in the training set. However, it is not applicable in unseen environments during inference since the model learns scene-specific features in the APR paradigm. ❷ RPR (in Fig. 3b), in contrast to APR, focuses on the relative features between image pairs during training. Although RPR has better generalizability compared to APR, it requires a large database during inference since RPR needs to retrieve a reference image similar to the query image to form an image pair as the model input. ❸ VO (in Fig. 3c) can generalize well in unseen environments but suffers from accumulated drift since it uses the previously predicted camera pose to predict the current pose. ❹ SPR (in Fig. 3d) is proposed to address the aforementioned problems. By choosing the first image in the sequence as the scene origin, SPR separates the coordinate system from particular scenes in the training set and identifies scene-agnostic features between the query and the origin image. This enables the model to perform more effectively in novel scenes without

the need for a large-scale database. It is worth noting that open trajectories in the SPR paradigm do not experience accumulated drift, as the regression of the current camera pose is not influenced by the preceding pose.

3.2. 360SPR: Established Dataset

Data Collection. We first collect pinhole images in size of 512×512 using the Habitat simulator [47, 55, 59] powered by HM3D [49] and Matterport3D [6] datasets and then stitch the pinhole images to obtain panoramas in 1024×2048. We use the same stitching tool as Matterport3D [6] dataset. For every sample point in the trajectories, we collect images with 3 elevations and 6 headings, resulting in 18 pinhole images with their corresponding camera poses. We select the camera pose of the 10-th pinhole image in the sequential sequence as the pose of the stitched panorama. To enable high-quality panoramic images, three inspectors manually checked all samples in the form of cross validation. The whole cleaning process took more than 300 hours. Fig. 1 showcases the data collection process. Within a navigable area of a scene, we randomly select two points as the starting and destination points. Then we calculate the shortest path between the two points using the Dijkstra [14] algorithm. Since the 360Loc [25] dataset doesn’t consider different sampling intervals and sensor heights, it’s difficult to satisfy the need for robust and accurate spatial awareness in various real-world applications. To this end, we sample trajectories in different lengths with varying sampling intervals between sampling points along the path. The trajectory length in 360SPR varies from 3m to 20m and the number of panoramas in one trajectory varies from 5 to 20. Moreover, three different robot heights with a sampling ratio of 1:1:2 are also taken into account, i.e., sweeping , quadruped , and humanoid  robots as shown in Fig. 1.

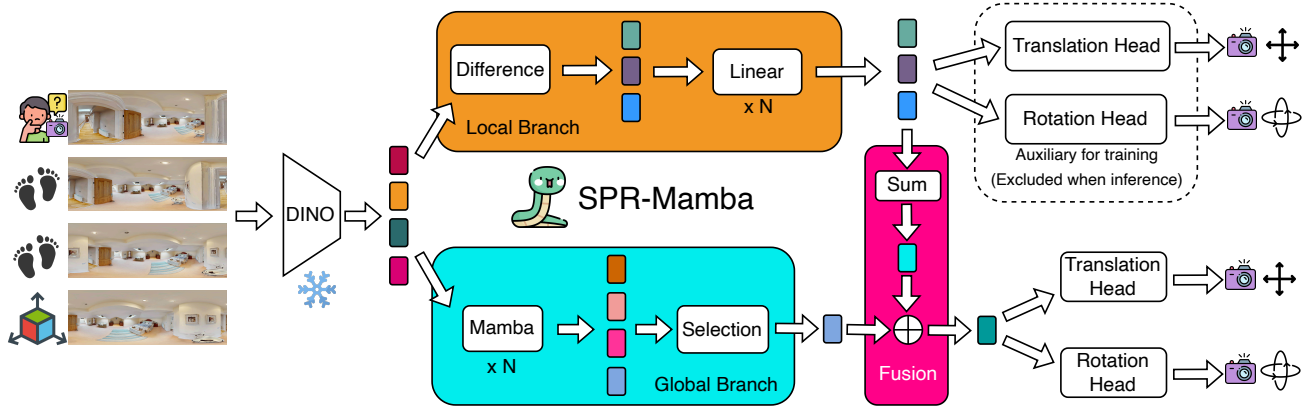


Figure 4. Architecture of SPR-Mamba. The local branch learns the relative poses between adjacent frames while the global branch focuses on the pose between the query and the first frame in the sequence. The auxiliary heads for the local branch are not necessary after training. Note that SPR-Mamba can handle arbitrary sequence length during inference and we use 4 images for the demonstration purpose.

Data Statistics. Table 1 lists the statistics of different visual localization datasets, including 7Scenes [19], 12Scenes [64], InLoc [60], Cambridge [30], NaVIP [75], LaMAR [54], and 360Loc [25]. All these datasets are photographic. As 360SPR is collected using the Habitat [47, 55, 59] simulator, the stitched panoramas are photorealistic rather than photographic, providing the potential for domain adaptation and generalization tasks. Besides, our 360SPR dataset focuses on indoor scenes with different sensor heights. Compared with other visual localization datasets, 360SPR showcases its value with more than 200K panoramas and 3.6M pinholes distributed in 20K trajectories among 270 scenes, addressing the data capacity and diversity problem of existing datasets for visual localization. For the visualization of the 360SPR panoramas, please refer to the supplementary.

3.3. SPR-Mamba Framework

Preliminary. State Space Models [21] (SSMs) and Mamba [13, 20] have been receiving increasing interest in the community due to their excellent performance in modeling sequential inputs and outputs. SSMs can be expressed by using a hidden state $h(t) \in \mathbb{R}^N$, parameter $\mathbf{A} \in \mathbb{R}^{N \times N}$, $\mathbf{B} \in \mathbb{R}^N$, $\mathbf{C} \in \mathbb{R}^N$ as follows:

$$h'(t) = \mathbf{A}h(t) + \mathbf{B}x(t), \quad (1)$$

$$y(t) = \mathbf{C}h(t). \quad (2)$$

Discretizing SSMs [21] by leveraging a timescale parameter Δ , Equations 1 and 2 can be reformulated as:

$$h_t = \bar{\mathbf{A}}h_{t-1} + \bar{\mathbf{B}}x_t, \quad (3)$$

$$y_t = \mathbf{C}h_t, \quad (4)$$

where

$$\bar{\mathbf{A}} = \exp(\Delta\mathbf{A}), \quad (5)$$

$$\bar{\mathbf{B}} = (\Delta\mathbf{A})^{-1}(\exp(\Delta\mathbf{A}) - \mathbf{I}) \cdot \Delta\mathbf{B}. \quad (6)$$

Compared to RNNs and Transformers [65], one advantage of SSMs is that they can be trained in parallel while doing inference in a sequential manner:

$$\bar{\mathbf{K}} = (\mathbf{C}\bar{\mathbf{B}}, \mathbf{C}\bar{\mathbf{A}}\bar{\mathbf{B}}, \dots, \mathbf{C}\bar{\mathbf{A}}^{L-1}\bar{\mathbf{B}}), \quad (7)$$

$$y = x * \bar{\mathbf{K}}. \quad (8)$$

Making SSMs input-sensitive, Mamba [13, 20] is proposed by modifying the following terms:

$$\bar{\mathbf{B}}_t = \text{Linear}_{\mathbf{B}}(x_t), \quad (9)$$














$$\bar{\mathbf{C}}_t = \text{Linear}_{\mathbf{C}}(x_t), \quad (10)$$

$$\Delta_t = \text{Softplus}(\text{Linear}_{\Delta}(x_t)). \quad (11)$$

SPR-Mamba. Given a sequence of images along a trajectory, SPR-Mamba is able to predict the camera pose of an arbitrary query image I_q w.r.t. I_1 for an arbitrary sequence length I_1, I_2, \dots, I_q . Fig. 4 presents the model architecture of the dual-branch SPR-Mamba and sets $q=4$ for the illustration purpose. SPR-Mamba utilizes a frozen DINO [5] to extract image features. The features are then fed into a local branch to attain frame-by-frame features and a global branch for the query-to-origin features. Within the local branch, SPR-Mamba first calculates the difference between two consecutive adjacent frames, *i.e.*, 4 image features result in 3 feature differences, followed by multiple linear layers. The auxiliary translation head then outputs 3 relative translations of the 3 frame-by-frame relative poses from 4 images. The same operation applies to the auxiliary rotation head during training. Note that the two auxiliary heads are not necessary after training.

In addition to the local branch focusing on the frame-by-frame camera poses, the 4 DINO-extracted features also go through the global branch for query-to-origin feature learning. The global branch is stacked by multiple Mamba blocks and finally selects the last hidden state from all outcomes outputted by the last Mamba block since the last one

Table 2. Comparison of different models using different paradigms in both **seen** and **unseen** environments on the **360SPR** dataset. The average median and average mean of Translation Error (TE in meters) and Rotation Error (RE in degrees) are reported.

Paradigm	Model	Source	Code	#Image	Average Median				Average Mean			
					TE (seen)	TE (unseen)	RE (seen)	RE (unseen)	TE (seen)	TE (unseen)	RE (seen)	RE (unseen)
APR	PoseNet [30]	ICCV	link	 ×1	10.12±0.3	30.25±0.2	10.22±0.3	47.15±0.2	10.13±0.2	29.54±0.2	10.23±0.2	46.02±0.1
	NeFeS [7]	CVPR	link	 ×1	3.29±0.3	27.88±0.9	3.40±0.3	49.01±0.8	3.20±0.3	29.04±1.1	3.22±0.2	47.51±0.9
	Marepo [8]	CVPR	link	 ×1	3.22±0.2	27.98±1.1	3.31±0.3	48.12±1.1	3.13±0.3	28.96±1.2	3.02±0.2	47.44±1.2
	VidLoc [11]	CVPR	link	 ×5	9.23±0.4	27.44±1.2	9.62±0.5	46.99±1.2	9.43±0.2	27.45±1.1	9.17±0.4	47.33±1.0
	MapNet [3]	CVPR	link	 ×5	9.23±0.2	27.12±1.2	9.45±0.6	47.22±1.3	9.45±0.2	26.71±1.2	9.34±0.3	46.15±0.7
	GL-Net [71]	CVPR	reimpl.	 ×5	8.61±0.2	27.45±1.0	9.31±0.4	47.01±0.4	8.91±0.4	28.44±1.1	8.89±0.4	48.21±1.0
RPR	NN-Net [35]	ICCVW	link	 ×1	10.93±0.3	12.84±0.4	10.32±0.2	22.65±0.2	10.34±0.3	12.89±0.2	10.25±0.5	22.22±0.4
	FAR [50]	CVPR	link	 ×1	10.06±0.3	11.85±0.3	9.51±0.3	21.04±0.2	10.02±0.3	11.22±0.4	10.26±0.4	21.16±0.5
	PanoPose [62]	CVPR	reimpl.	 ×1	10.01±0.4	10.91±0.3	9.02±0.4	20.01±0.3	10.23±0.2	11.03±0.3	10.12±0.4	20.55±0.4
VO	DPVO [61]	NeurIPS	link	 ×5	3.88±0.3	4.02±0.4	4.12±0.2	4.38±0.4	3.71±0.2	3.92±0.3	4.35±0.3	4.44±0.4
	LEAP-VO [9]	CVPR	link	 ×5	3.77±0.4	3.89±0.3	4.22±0.2	4.30±0.4	3.72±0.3	3.85±0.3	4.30±0.4	4.33±0.2
	XVO [34]	ICCV	link	 ×5	4.11±0.3	4.25±0.3	4.02±0.3	4.21±0.3	3.68±0.4	3.88±0.4	4.22±0.2	4.27±0.4
SPR	SPR-Mamba (ours)	CVPR	link	 ×5	3.32±0.3	3.85±0.3	3.43±0.3	3.97±0.4	3.22±0.2	3.78±0.4	3.31±0.3	3.91±0.3

aggregates all information from I_1 to I_q . The selected hidden state is then fused with the features outputted by the local branch. Like the local branch, a translation and a rotation head take the fused feature as input and output the relative query-to-origin camera translation and rotation, respectively. Since Mamba is capable of doing inference sequentially, SPR-Mamba can continuously handle the upcoming panoramas during inference with linear complexity.

4. Experiments

4.1. Datasets

360SPR. Our new dataset provides over 200K photorealistic panoramas and the corresponding camera poses in 20K trajectories distributed in 270 scenarios. All panoramas have 1024×2048 resolution captured at 3 sensor heights (sweeping 🌀, quadruped 🦘, and humanoid 🤖 robots).

360Loc. The dataset [25] consists of 9.3K photographic panoramas with corresponding camera poses in 18 independent trajectories distributed in 4 scenarios, namely Concourse, Hall, Atrium, and Piatium. All panoramas have a 3072×6144 resolution captured at a fixed sensor height.

4.2. Implementation details

We train the SPR-Mamba model from scratch without any pretraining except for a frozen DINO [5] as the feature extractor. The SPR-Mamba is trained with an A100 GPU for 150 epochs. The AdamW [41] optimizer is applied with an initial learning rate of $1e^{-4}$. The training is warmed by a linear scheduler for the first 10 epochs followed by a cosine annealing strategy. To facilitate the training and inference, we resize the panoramic images to 320×640 for the 360SPR and 392×770 for the 360Loc [25] dataset. SPR-Mamba is trained with a sequence length of 5 images and uses the last one as the query image. Applying a batch size of 8 results in 40 images within a batch. We use L1 loss to supervise

every camera pose:

$$L = \alpha \|\hat{\mathbf{t}} - \mathbf{t}\|_1 + \beta \|\hat{\mathbf{q}} - \mathbf{q}\|_1, \quad (12)$$

where $\hat{\mathbf{t}}$ and $\hat{\mathbf{q}}$ are the SPR-Mamba translation and rotation output, \mathbf{t} and \mathbf{q} are the translation and rotation ground-truth, respectively. α and β are two scaling factors balancing the translation and rotation losses. Empirically, we choose $\alpha = 1$ and $\beta = 10$. Following [63], we parameterize camera rotation as the logarithm of a unit quaternion, equivalent to the axis-angle representation up to scale [23]. This operation avoids the need for additional constraints to ensure a valid rotation. Logarithmic mapping of the unit quaternion $\mathbf{q} = [u, \mathbf{v}]$ is done via:

$$\log(\mathbf{q}) = \begin{cases} \frac{\mathbf{v}}{\|\mathbf{v}\|_2} \cos^{-1}(u) & \text{if } \|\mathbf{v}\|_2 \neq 0, \\ \mathbf{0} & \text{otherwise,} \end{cases} \quad (13)$$














where u and \mathbf{v} are the real and the imaginary part of a unit quaternion, respectively. The logarithmic form $\mathbf{w} = \log(\mathbf{q})$ can be converted back to a unit quaternion by the exponential mapping, $\exp(\mathbf{w}) = [\cos(\|\mathbf{w}\|_2), \frac{\mathbf{w}}{\|\mathbf{w}\|_2} \sin(\|\mathbf{w}\|_2)]$.

4.3. Quantitative Results

To benchmark the proposed SPR task, we conduct experiments on two panoramic datasets, *i.e.*, our 360SPR and the 360Loc [25] dataset. All models are trained three times for a fair comparison. The average median and mean of translation error (TE in meters) and rotation error (RE in degrees) among all scenes with uncertainty are reported.

Results on 360SPR. Table 2 shows the comparison results of different paradigms. In the **unseen** setting, we leave out 15 of 270 scenes as unknown during training and use them for testing. Image-based APR [7, 8, 30] and sequence-based APR [3, 11, 71] models have much lower generalizability in unseen environments. For example, Marepo [8] obtains the average median of TE/RE in 27m/48° even though it

Table 3. Comparison of different models using different paradigms in both **seen** and **unseen** environments on the **360Loc** dataset. The average median and average mean of Translation Error (TE in meters) and Rotation Error (RE in degrees) are reported.

Paradigm	Model	Source	Code	#Image	Average Median				Average Mean			
					TE (seen)	TE (unseen)	RE (seen)	RE (unseen)	TE (seen)	TE (unseen)	RE (seen)	RE (unseen)
APR	PoseNet [30]	ICCV	link	 ×1	8.23±0.2	28.55±1.4	8.34±0.3	45.12±1.1	8.26±0.2	27.32±1.6	8.55±0.3	46.02±1.3
	NeFeS [7]	CVPR	link	 ×1	1.27 ±0.2	25.78±0.7	1.12±0.3	47.12±0.7	1.21±0.3	27.12±1.0	1.12±0.2	45.63±0.9
	Marepo [8]	CVPR	link	 ×1	1.31±0.2	25.75±1.2	1.11 ±0.3	46.21±1.1	1.17 ±0.2	26.81±1.4	1.03 ±0.2	45.11±1.5
	VidLoc [11]	CVPR	link	 ×5	7.36±0.4	25.33±1.1	7.72±0.5	44.39±1.1	7.66±0.3	25.13±1.0	7.17±0.4	45.19±1.3
	MapNet [3]	CVPR	link	 ×5	7.03±0.4	25.21±1.3	7.32±0.7	45.43±1.3	7.31±0.4	24.66±1.1	7.23±0.3	44.51±1.1
	GL-Net [71]	CVPR	reimpl.	 ×5	6.45±0.3	25.72±1.3	7.44±0.4	45.23±1.1	6.94±0.4	26.65±1.2	7.01±0.3	46.11±1.0
RPR	NN-Net [35]	ICCVW	link	 ×1	8.98±0.2	10.97±0.3	8.13±0.4	20.51±0.2	8.53±0.4	10.77±0.4	8.37±0.3	20.01±0.4
	FAR [50]	CVPR	link	 ×1	7.99±0.3	9.35±0.3	7.42±0.4	19.01±0.2	8.02±0.3	9.13±0.3	8.16±0.4	19.05±0.4
	PanoPose [62]	CVPR	reimpl.	 ×1	7.81±0.4	9.12±0.4	7.35±0.4	18.94±0.4	7.99±0.3	9.01±0.3	8.06±0.3	19.00±0.5
VO	DPVO [61]	NeurIPS	link	 ×5	2.22±0.3	2.40±0.3	1.72±0.3	1.84±0.2	2.30±0.3	2.43±0.3	1.82±0.3	2.01±0.3
	LEAP-VO [9]	CVPR	link	 ×5	2.54±0.2	2.71±0.3	1.68±0.2	2.03±0.3	2.46±0.3	2.66±0.3	1.83±0.3	2.01±0.3
	XVO [34]	ICCV	link	 ×5	2.43±0.3	2.56±0.2	1.80±0.4	1.99±0.3	2.51±0.4	2.69±0.3	1.68±0.2	1.92±0.3
SPR	SPR-Mamba (ours)	CVPR	link	 ×5	1.43±0.3	1.94 ±0.3	1.21±0.2	1.44 ±0.2	1.23±0.3	1.87 ±0.3	1.17±0.3	1.28 ±0.2

achieves the best performance in seen environments. The reason for the poor generalization is that models in the APR paradigm learn scene-specific features to predict the absolute camera pose. Therefore, they fail to predict poses in novel scenes. On the other hand, RPR models generalize better than APR models since they focus on the relative features between frames which are less related to the scenes. VO models can generalize well in unseen environments. However, compared to SPR, VO performs worse due to the unavoidable accumulated drift. Our SPR-Mamba in the SPR paradigm significantly outperforms all APR, RPR and VO models in unknown environments, yielding the best performance (the average median of TE/RE in **3.85m/3.97°**). The error reduction of TE/RE are over **23m/43°**↓ compared to APR, and **7m/16°**↓ compared to RPR. The SPR paradigm disentangles the coordinate system from specific scenes and predicts the relative pose between the query and origin frame by learning the scene-agnostic features along trajectories. These results show that our SPR paradigm and model generalize better in unknown environments. In the **seen** setting, 20% trajectories of training scenes are excluded for testing. Surprisingly, the performance of our SPR-Mamba is on par with the best image-based APR method Marepo [8].

Results on 360Loc. Table 3 shows the comparison on the 360Loc [25] dataset. In the **unseen** environments, we perform cross-validation, in which 1 of 4 scenes is excluded as unknown and for testing, while the remaining 3 scenes for training. The same operation repeats until all 4 scenes are tested. We then average the results of 4 scenes. Consistent with the results on 360SPR, all APR models generalize poorly in unseen scenes on the 360Loc [25] dataset. RPR models showcase the superiority in unseen environments compared to both image-based and sequence-based APR models. Although VO models perform better than APR and RPR models in unknown environments,




Figure 5. Comparison of VO and SPR paradigms for unknown scenes with different sequence lengths of 360Loc [25] dataset.

they underperform models in the SPR paradigm. Our SPR-Mamba outperforms all models, obtaining the best performance (the average median of TE/RE in **1.94m/1.44°**) with over **7m/17°**↓ error reduction of the RPR paradigm. In the **seen** testing, the data splitting follows the official 360Loc [25] dataset. The performance gap between SPR-Mamba and Marepo [8] is marginal with 0.12m TE and 0.1° RE. The remarkable performance on the real-world panoramic dataset proves the effectiveness of our SPR-Mamba in both seen and unseen environments.

4.4. Ablation Study

Comparison with VO. Fig. 5 shows the comparison between VO and SPR paradigms. Compared with VO, SPR predicts the camera pose relative to the first origin frame instead of the previous frame while VO predicts the relative pose between two consecutive adjacent frames. Since VO predicts the current pose based on the previous one, the accumulated drift in open trajectories is unavoidable. In contrast, our SPR paradigm directly estimates the pose according to the origin, avoiding the need for loop closure detection to eliminate the drift. We train TSformer-VO [17] in both paradigms on the 360Loc [25] dataset with different sequence lengths, from 5 to 20 images. The model is trained with 3 scenes and evaluated with the last remaining scene in

Table 4. Ablation study of SPR-Mamba components on 360SPR dataset. Translation/Rotation Errors (TE/RE in $m/^{\circ}$) are reported.

Model	Aux. Heads	Local Branch	Global Branch	TE	RE
① SPR-Mamba	✓	✓	✗	10.33 \pm 1.3	10.68 \pm 1.4
② SPR-Mamba	✗	✗	✓	4.72 \pm 0.7	5.11 \pm 0.5
③ SPR-Mamba	✗	✓	✓	4.32 \pm 0.6	4.67 \pm 0.5
④ SPR-Mamba	✓	✓	✓	3.85 \pm 0.3	3.97 \pm 0.4

a cross-validation manner. We train the model three times and report the average median of TE and RE with standard deviation in shadow in Fig. 5. It can be observed that the TE and RE of TSformer-VO [17] become larger as the sequence becomes longer in both SPR and VO settings. However, the performance gap widens as the sequence length increases since the SPR paradigm does not experience drift, whereas the VO paradigm does. The result confirms that our SPR paradigm can better handle long sequences.

Component Ablation. To further analyze the effectiveness of different components of the proposed SPR-Mamba, we ablate the local branch, global branch, and auxiliary training heads. The average median of TE and RE in unknown environments on the 360SPR dataset are reported in Table 4. ① Without the global branch, SPR-Mamba degenerates into a simple model consisting of several linear layers, predicting the consecutive adjacent frames like VO. Apart from the aforementioned variant, ② SPR-Mamba without the local branch and the auxiliary regression heads cannot achieve satisfying results, either. The model only focuses on the query-to-origin relative features, neglecting the local relative features which help improve overall model performance. Compared with the previous variant, this model has a performance gain over $5m/5^{\circ}$ for TE and RE, respectively. ③ SPR-Mamba with a local and global branch but without the auxiliary regression heads, achieves even better performance since the model is able to capture the local relative features even though there is no supervision for the relative poses of the consecutive adjacent frames. ④ SPR-Mamba with all components, performs the best. In addition to the local and global feature learning, the auxiliary heads provide more supervision via their loss functions for the local branch, which further benefits the global feature learning through gradient and backpropagation.

Cross-data Evaluation. We conduct the cross-evaluation experiment between 360Loc [25] and 360SPR datasets. Table 5a demonstrates the average median of TE and RE for unseen environments of both datasets. When evaluated on 360Loc [25], SPR-Mamba trained on 360SPR achieves comparable results to those trained on 360Loc [25]. The performance gap of $0.2m/0.23^{\circ}$ on TE/RE is neglectable. In contrast, SPR-Mamba trained on 360Loc [25] is not able to perform well on the 360SPR dataset because 360SPR is more diverse and challenging. It’s difficult for a model to generalize well when trained on a dataset with a single sensor height, similar scenes and limited panoramic data.

Table 5. Evaluation between datasets and between sensor heights. Translation/Rotation Errors (TE/RE in $m/^{\circ}$) are reported.

		Evaluate				
Train	Evaluate	360Loc	360SPR			
360Loc		1.94 / 1.44	6.02 / 6.11	3.88 / 3.97	5.33 / 5.41	5.45 / 5.57
		5.23 / 5.30	3.93 / 4.01	5.23 / 5.30	5.44 / 5.55	4.90 / 5.10
		5.18 / 5.29	5.21 / 5.33	5.18 / 5.29	3.78 / 3.86	4.49 / 4.59
360SPR		2.14 / 1.67	3.85 / 3.97	3.91 / 4.03	3.90 / 3.99	3.79 / 3.92
				3.91 / 4.03	3.90 / 3.99	3.85 / 3.97

(a) Evaluation cross datasets.

(b) Evaluation cross sensor heights.

The performance gap, in this case, reaches $2.17m/2.14^{\circ}$ on TE/RE. This ablation study shows the value of creating a large-scale panoramic dataset like 360SPR with over 200K panoramas and 270 scenes at 3 different sensor heights, making a model more robust to real-world cases.

Cross-sensor Evaluation. Apart from the cross-dataset evaluation, we also ablate the influence of different sensor heights in unknown environments. As shown in Table 5b, we train SPR-Mamba model on 4 splits of 360SPR, namely the data only at $0.1m$, data only at $0.5m$, data only at $1.7m$, and a mixture of all heights . It can be observed that the domain gap exists across different heights, e.g., SPR-Mamba trained on pure $1.7m$ data achieves a $3.78m$ average median of translation error evaluated at the same height, while the error increases by $1.4m$ \uparrow when evaluated at the pure $0.1m$ data. Compared with the training at a single height, SPR-Mamba showcases excellent robustness when trained on the complete 360SPR dataset, achieving the best/second-best results in all splits. This ablation study proves the necessity of creating a dataset with different sensor heights for accurate and robust camera pose regression.

5. Conclusion

We propose Scene-agnostic Pose Regression (SPR) to enhance APR generalization and reduce RPR’s database dependency. To explore SPR, we create a large-scale panoramic dataset (360SPR) with over 200K photorealistic panoramic images and $3.6M$ pinhole images in 270 scenes. To address SPR, a dual-branch SPR-Mamba model is constructed with an SSM-based mechanism, which showcases its superiority in the camera pose regression task in unseen environments with a $7m/16^{\circ}$ error degradation in translation and rotation compared with APR and RPR models. We hope this work has the potential to advance the field of camera pose regression, providing generalized visual localization in unknown scenes.

Limitations and Future Work. Although the SPR paradigm is capable of predicting camera poses in unknown environments, the poses are relative to the origin frame. No information about the absolute poses is available. Besides, the ability of SPR-Mamba to handle image distortions in panoramas will be further improved in our future work.

Acknowledgments

This work was supported in part by the Ministry of Science, Research and the Arts of Baden-Württemberg (MWK) through the Cooperative Graduate School Accessibility through AI-based Assistive Technology (KATE) under Grant BW6-03, in part by Karlsruhe House of Young Scientists (KHYS), in part by the Helmholtz Association Initiative and Networking Fund on the HAICORE@KIT and HOREKA@KIT partition, in part by the National Natural Science Foundation of China (No. 62473139), and in part by the National Key RD Program under Grant 2022YFB4701400.

References

- [1] Hernán Badino, Akihiro Yamamoto, and Takeo Kanade. Visual odometry by multi-frame feature integration. In *ICCV*, 2013. 3
- [2] Vassileios Balntas, Shuda Li, and Victor Prisacariu. RelocNet: Continuous metric learning relocalisation using neural nets. In *ECCV*, 2018. 3
- [3] Samarth Brahmhatt, Jinwei Gu, Kihwan Kim, James Hays, and Jan Kautz. Geometry-aware learning of maps for camera localization. In *CVPR*, 2018. 3, 6, 7
- [4] Ke Cao, Ruiping Liu, Ze Wang, Kunyu Peng, Jiaming Zhang, Junwei Zheng, Zhifeng Teng, Kailun Yang, and Rainer Stiefelhofen. Tightly-coupled LiDAR-visual SLAM based on geometric features for mobile agents. In *ROBIO*, 2023. 3
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 5, 6
- [6] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3D: Learning from RGB-D data in indoor environments. *3DV*, 2017. 2, 4
- [7] Shuai Chen, Yash Bhalgat, Xinghui Li, Jia-Wang Bian, Kejie Li, Zirui Wang, and Victor Adrian Prisacariu. Neural refinement for absolute pose regression with feature synthesis. In *CVPR*, 2024. 6, 7
- [8] Shuai Chen, Tommaso Cavallari, Victor Adrian Prisacariu, and Eric Brachmann. Map-relative pose regression for visual re-localization. In *CVPR*, 2024. 3, 6, 7
- [9] Weirong Chen, Le Chen, Rui Wang, and Marc Pollefeys. LEAP-VO: Long-term effective any point tracking for visual odometry. In *CVPR*, 2024. 6, 7
- [10] Hsiang-Jen Chien, Chen-Chi Chuang, Chia-Yen Chen, and Reinhard Klette. When to use what feature? SIFT, SURF, ORB, or A-KAZE features for monocular visual odometry. In *IVCNZ*, 2016. 3
- [11] Ronald Clark, Sen Wang, Andrew Markham, Niki Trigoni, and Hongkai Wen. VidLoc: A deep spatio-temporal model for 6-DoF video-clip relocalization. In *CVPR*, 2017. 3, 6, 7
- [12] Jiajia Dai, Xiaoxi Gong, Yida Li, Jun Wang, and Mingqiang Wei. Self-supervised deep visual odometry based on geometric attention model. *IEEE Transactions on Intelligent Transportation Systems*, 2023. 3
- [13] Tri Dao and Albert Gu. Transformers are SSMs: Generalized models and efficient algorithms through structured state space duality. In *ICML*, 2024. 5
- [14] Edsger W. Dijkstra. A note on two problems in connexion with graphs. In *Edsger Wybe Dijkstra: his life, work, and legacy*. 2022. 4
- [15] Mingyu Ding, Zhe Wang, Jiankai Sun, Jianping Shi, and Ping Luo. CamNet: Coarse-to-fine retrieval for camera relocalization. In *ICCV*, 2019. 3
- [16] Jakob Engel, Jürgen Sturm, and Daniel Cremers. Semi-dense visual odometry for a monocular camera. In *ICCV*, 2013. 3
- [17] André O. Françani and Marcos R. O. A. Máximo. Transformer-based model for monocular visual odometry: A video understanding approach. *arXiv preprint arXiv:2305.06121*, 2023. 7, 8
- [18] Mariia Gladkova, Rui Wang, Niclas Zeller, and Daniel Cremers. Tight integration of feature-based relocalization in monocular direct visual odometry. In *ICRA*, 2021. 3
- [19] Ben Glocker, Shahram Izadi, Jamie Shotton, and Antonio Criminisi. Real-time RGB-D camera relocalization. In *ISMAR*, 2013. 4, 5
- [20] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023. 5
- [21] Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*, 2021. 5
- [22] Banglei Guan, Ji Zhao, Zhang Li, Fang Sun, and Friedrich Fraundorfer. Relative pose estimation with a single affine correspondence. *IEEE Transactions on Cybernetics*, 2022. 3
- [23] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 6
- [24] Jie Hu, Junwei Zheng, Jiale Wei, Jiaming Zhang, and Rainer Stiefelhofen. Deformable mamba for wide field of view segmentation. *arXiv preprint arXiv:2411.16481*, 2024. 2
- [25] Huajian Huang, Changkun Liu, Yipeng Zhu, Hui Cheng, Tristan Braud, and Sai-Kit Yeung. 360Loc: A dataset and benchmark for omnidirectional visual localization with cross-device queries. In *CVPR*, 2024. 2, 4, 5, 6, 7, 8
- [26] Gazi Karam Illahi, Ashutosh Vaishnav, Teemu Kämäräinen, Matti Siekkinen, and Mario Di Francesco. Learning to predict head pose in remotely-rendered virtual reality. In *MM-Sys*, 2023. 1
- [27] Zhong Ji, Keqin Nan, Yan Xu, Haoyuan Wang, Xuening Li, and Jiale Bai. Global-context-aware visual odometry system with epipolar-geometry-constrained loss function. *IEEE Transactions on Instrumentation and Measurement*, 2024. 3
- [28] Alex Kendall and Roberto Cipolla. Modelling uncertainty in deep learning for camera relocalization. In *ICRA*, 2016. 3
- [29] Alex Kendall and Roberto Cipolla. Geometric loss functions for camera pose regression with deep learning. In *CVPR*, 2017.
- [30] Alex Kendall, Matthew Grimes, and Roberto Cipolla. PoseNet: A convolutional network for real-time 6-DOF camera relocalization. In *ICCV*, 2015. 2, 3, 4, 5, 6, 7

- [31] Nicola Krombach, David Droeschel, and Sven Behnke. Combining feature-based and direct methods for semi-dense real-time stereo visual odometry. In *IAS*, 2016. 3
- [32] Nicola Krombach, David Droeschel, Sebastian Houben, and Sven Behnke. Feature-based visual odometry prior for real-time semi-dense stereo SLAM. *Robotics and Autonomous Systems*, 2018. 3
- [33] Xin-Yu Kuo, Chien Liu, Kai-Chen Lin, and Chun-Yi Lee. Dynamic attention-based visual odometry. In *CVPRW*, 2020. 3
- [34] Lei Lai, Zhongkai Shangguan, Jimuyang Zhang, and Eshed Ohn-Bar. XVO: Generalized visual odometry via cross-modal self-training. In *ICCV*, 2023. 6, 7
- [35] Zakaria Laskar, Iaroslav Melekhov, Surya Kalia, and Juho Kannala. Camera relocalization by computing pairwise relative poses using convolutional neural network. In *ICCVW*, 2017. 3, 6, 7
- [36] Haoang Li, Ji Zhao, Jean-Charles Bazin, Wen Chen, Kai Chen, and Yun-Hui Liu. Line-based absolute and relative camera pose estimation in structured environments. In *IROS*, 2019. 3
- [37] Shunkai Li, Xin Wang, Yingdian Cao, Fei Xue, Zike Yan, and Hongbin Zha. Self-supervised deep visual odometry with online adaptation. In *CVPR*, 2020. 3
- [38] Shunkai Li, Xin Wu, Yingdian Cao, and Hongbin Zha. Generalizing to the open world: Deep visual odometry with online adaptation. In *CVPR*, 2021. 3
- [39] Jingyu Lin, Jiaqi Gu, Bojian Wu, Lubin Fan, Renjie Chen, Ligang Liu, and Jieping Ye. Learning neural volumetric pose features for camera localization. In *ECCV*, 2024. 3
- [40] Ruyi Liu, Zhengzhe Liu, Haoyu Zhang, Guodao Zhang, Jianhua Zhang, Sunbo, Weiguo Sheng, Xiufeng Liu, and Yaochu Jin. ColVO: Colonoscopic visual odometry considering geometric and photometric consistency. In *MM*, 2024. 3
- [41] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. 6
- [42] Nico Messikommer, Giovanni Cioffi, Mathias Gehrig, and Davide Scaramuzza. Reinforcement learning meets visual odometry. In *ECCV*, 2024. 3
- [43] Arthur Moreau, Nathan Piasco, Dzmitry Tsishkou, Bogdan Stanculescu, and Arnaud de La Fortelle. CoordiNet: Uncertainty-aware pose regressor for reliable vehicle localization. In *WACV*, 2022. 1
- [44] Luca Morelli, Francesco Ioli, Raniero Beber, Fabio Menna, Fabio Remondino, and Alfonso Vitti. COLMAP-SLAM: A framework for visual odometry. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2023. 3
- [45] Lia Morra, Andrea Biondo, Nicola Poerio, and Fabrizio Lamberti. MIXO: Mixture of experts-based visual odometry for multicamera autonomous systems. *IEEE Transactions on Consumer Electronics*, 2023. 3
- [46] Jens Naumann, Binbin Xu, Stefan Leutenegger, and Xingxing Zuo. NeRF-VO: Real-time sparse visual odometry with neural radiance fields. *IEEE Robotics and Automation Letters*, 2024. 3
- [47] Xavier Puig, Eric Undersander, Andrew Szot, Mikael Dal-laire Cote, Tsung-Yen Yang, Ruslan Partsey, Ruta Desai, Alexander William Clegg, Michal Hlavac, So Yeon Min, et al. Habitat 3.0: A co-habitat for humans, avatars and robots. In *ICLR*, 2024. 2, 4, 5
- [48] Chengyu Qiao, Zhiyu Xiang, Yuangang Fan, Tingming Bai, Xijun Zhao, and Jingyun Fu. TransAPR: Absolute camera pose regression with spatial and temporal attention. *IEEE Robotics and Automation Letters*, 2023. 1
- [49] Santhosh Kumar Ramakrishnan, Aaron Gokaslan, Erik Wijmans, Oleksandr Maksymets, Alexander Clegg, John M Turner, Eric Undersander, Wojciech Galuba, Andrew Westbury, Angel X Chang, Manolis Savva, Yili Zhao, and Dhruv Batra. Habitat-Matterport 3D Dataset (HM3D): 1000 Large-scale 3D environments for embodied AI. In *NeurIPS*, 2021. 2, 4
- [50] Chris Rockwell, Nilesh Kulkarni, Linyi Jin, Jeong Joon Park, Justin Johnson, and David F. Fouhey. FAR: Flexible accurate and robust 6DoF relative camera pose estimation. In *CVPR*, 2024. 6, 7
- [51] Jiahao Ruan, Li He, Yisheng Guan, and Hong Zhang. Combining scene coordinate regression and absolute pose regression for visual relocalization. In *ICRA*, 2023. 1
- [52] Soham Saha, Girish Varma, and C. V. Jawahar. Improved visual relocalization by discovering anchor points. In *BMVC*, 2018. 3
- [53] Paul-Edouard Sarlin, Ajaykumar Unagar, Måns Larsson, Hugo Germain, Carl Toft, Viktor Larsson, Marc Pollefeys, Vincent Lepetit, Lars Hammarstrand, Fredrik Kahl, and Torsten Sattler. Back to the feature: Learning robust camera localization from pixels to pose. In *CVPR*, 2021. 3
- [54] 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 *ECCV*, 2022. 1, 4, 5
- [55] Manolis Savva, Jitendra Malik, Devi Parikh, Dhruv Batra, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, and Vladlen Koltun. Habitat: A platform for embodied AI research. In *ICCV*, 2019. 2, 4, 5
- [56] Yoli Shavit, Ron Ferens, and Yosi Keller. Learning multi-scene absolute pose regression with transformers. In *ICCV*, 2021. 3
- [57] Shihao Shen, Yilin Cai, Wenshan Wang, and Sebastian Scherer. DytanVO: Joint refinement of visual odometry and motion segmentation in dynamic environments. In *ICRA*, 2023. 3
- [58] Libo Sun, Wei Yin, Enze Xie, Zhengrong Li, Changming Sun, and Chunhua Shen. Improving monocular visual odometry using learned depth. *IEEE Transactions on Robotics*, 2022. 3
- [59] Andrew Szot, Alexander Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John M. Turner, Noah Maestre, Mustafa Mukadam, Devendra Singh Chaplot, Oleksandr Maksymets, Aaron Gokaslan, Vladimir Vondrus, Sameer Dharur, Franziska Meier, Wojciech Galuba, Angel X. Chang,

- Zsolt Kira, Vladlen Koltun, Jitendra Malik, Manolis Savva, and Dhruv Batra. Habitat 2.0: Training home assistants to rearrange their habitat. In *NeurIPS*, 2021. [2](#), [4](#), [5](#)
- [60] 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 *CVPR*, 2018. [4](#), [5](#)
- [61] Zachary Teed, Lahav Lipson, and Jia Deng. Deep patch visual odometry. In *NeurIPS*, 2023. [3](#), [6](#), [7](#)
- [62] Diantao Tu, Hainan Cui, Xianwei Zheng, and Shuhan Shen. PanoPose: Self-supervised relative pose estimation for panoramic images. In *CVPR*, 2024. [6](#), [7](#)
- [63] Mehmet Ozgur Turkoglu, Eric Brachmann, Konrad Schindler, Gabriel J. Brostow, and Áron Monszpart. Visual camera re-localization using graph neural networks and relative pose supervision. In *3DV*, 2021. [2](#), [3](#), [6](#)
- [64] Julien P. C. Valentin, Angela Dai, Matthias Nießner, Pushmeet Kohli, Philip H. S. Torr, Shahram Izadi, and Cem Keskin. Learning to navigate the energy landscape. In *3DV*, 2016. [4](#), [5](#)
- [65] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. [3](#), [5](#)
- [66] Brandon Wagstaff, Valentin Peretroukhin, and Jonathan Kelly. Self-supervised deep pose corrections for robust visual odometry. In *ICRA*, 2020. [3](#)
- [67] Sen Wang, Ronald Clark, Hongkai Wen, and Niki Trigoni. DeepVO: Towards end-to-end visual odometry with deep recurrent convolutional neural networks. In *ICRA*, 2017. [3](#)
- [68] Jiale Wei, Junwei Zheng, Ruiping Liu, Jie Hu, Jiaming Zhang, and Rainer Stiefelhagen. OneBEV: Using one panoramic image for bird’s-eye-view semantic mapping. In *ACCV*, 2024. [2](#)
- [69] Jian Wu, Liwei Ma, and Xiaolin Hu. Delving deeper into convolutional neural networks for camera relocalization. In *ICRA*, 2017. [3](#)
- [70] Fei Xue, Xin Wang, Zike Yan, Qiuyuan Wang, Junqiu Wang, and Hongbin Zha. Local supports global: Deep camera relocalization with sequence enhancement. In *ICCV*, 2019. [3](#)
- [71] Fei Xue, Xin Wu, Shaojun Cai, and Junqiu Wang. Learning multi-view camera relocalization with graph neural networks. In *CVPR*, 2020. [3](#), [6](#), [7](#)
- [72] Chenhao Yang, Yuyi Liu, and Andreas Zell. RCPNet: Deep-learning based relative camera pose estimation for UAVs. In *ICUAS*, 2020. [3](#)
- [73] Nan Yang, Lukas von Stumberg, Rui Wang, and Daniel Cremers. D3VO: Deep depth, deep pose and deep uncertainty for monocular visual odometry. In *CVPR*, 2020. [3](#)
- [74] Zhenpei Yang, Jeffrey Z. Pan, Linjie Luo, Xiaowei Zhou, Kristen Grauman, and Qixing Huang. Extreme relative pose estimation for RGB-D scans via scene completion. In *CVPR*, 2019. [3](#)
- [75] Jun Yu, Yifan Zhang, Badrinadh Aila, and Vinod Nambodiri. NaVIP: An image-centric indoor navigation solution for visually impaired people. *arXiv preprint arXiv:2410.18109*, 2024. [4](#), [5](#)
- [76] Leijian Yu, Erfu Yang, Beiya Yang, Zixiang Fei, and Cong Niu. A robust learned feature-based visual odometry system for UAV pose estimation in challenging indoor environments. *IEEE Transactions on Instrumentation and Measurement*, 2023. [3](#)
- [77] Saiping Zhang, Dongyang Jin, Yuchao Dai, and Fuzheng Yang. Relative pose estimation for light field cameras based on LF-point-LF-point correspondence model. *IEEE Transactions on Image Processing*, 2022. [3](#)
- [78] Sen Zhang, Jing Zhang, and Dacheng Tao. Towards scale consistent monocular visual odometry by learning from the virtual world. In *ICRA*, 2022. [3](#)
- [79] Junwei Zheng, Ruiping Liu, Yufan Chen, Kunyu Peng, Chengzhi Wu, Kailun Yang, Jiaming Zhang, and Rainer Stiefelhagen. Open panoramic segmentation. In *European Conference on Computer Vision (ECCV)*, 2024. [2](#)
- [80] Qunjie Zhou, Torsten Sattler, Marc Pollefeys, and Laura Leal-Taixe. To learn or not to learn: Visual localization from essential matrices. In *ICRA*, 2020. [3](#)
- [81] Yuliang Zou, Pan Ji, Quoc-Huy Tran, Jia-Bin Huang, and Manmohan Chandraker. Learning monocular visual odometry via self-supervised long-term modeling. In *ECCV*, 2020. [3](#)