

ConDo: Continual Domain Expansion for Absolute Pose Regression

Zijun Li^{1*}, Zhipeng Cai^{2*†}, Bochun Yang¹, Xuelun Shen¹,
Siqi Shen¹, Xiaoliang Fan¹, Michael Paulitsch², Cheng Wang^{1†}

¹Fujian Key Laboratory of Sensing and Computing for Smart Cities, Xiamen University, China

²Intel Labs

{lizijun;yangbc;xuelun}@stu.xmu.edu.cn, {zhipeng.cai;michael.paulitsch}@intel.com,

{siqishen;fanxiaoliang;cwang}@xmu.edu.cn

Abstract

Visual localization is a fundamental machine learning problem. Absolute Pose Regression (APR) trains a scene-dependent model to efficiently map an input image to the camera pose in a pre-defined scene. However, many applications have continually changing environments, where inference data at novel poses or scene conditions (weather, geometry) appear after deployment. Training APR on a fixed dataset leads to overfitting, making it fail catastrophically on challenging novel data. This work proposes Continual Domain Expansion (ConDo), which continually collects unlabeled inference data to update the deployed APR. Instead of applying standard unsupervised domain adaptation methods which are ineffective for APR, ConDo effectively learns from unlabeled data by distilling knowledge from scene-agnostic localization methods. By sampling data uniformly from historical and newly collected data, ConDo can effectively expand the generalization domain of APR. Large-scale benchmarks with various scene types are constructed to evaluate models under practical (long-term) data changes. ConDo consistently and significantly outperforms baselines across architectures, scene types, and data changes. On challenging scenes (Fig. 1), it reduces the localization error by $> 7x$ (14.8m vs 1.7m). Analysis shows the robustness of ConDo against compute budgets, replay buffer sizes and teacher prediction noise. Comparing to model re-training, ConDo achieves similar performance up to 25x faster.

Code — <https://github.com/ZijunLi7/ConDo>

1 Introduction

Localizing an image in a given scene is a fundamental machine learning problem. The scene is defined by a set of *reference* images with known camera poses, and the task is to return the camera pose of a query image.

Different types of methods have been developed for visual localization. Retrieval-based methods search for reference images similar to the input and use their poses as the output (Torii et al. 2015; Arandjelovic et al. 2016). These methods require storing the reference images during inference, which introduces memory overheads. There are also methods

*These authors contributed equally.

†Corresponding authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

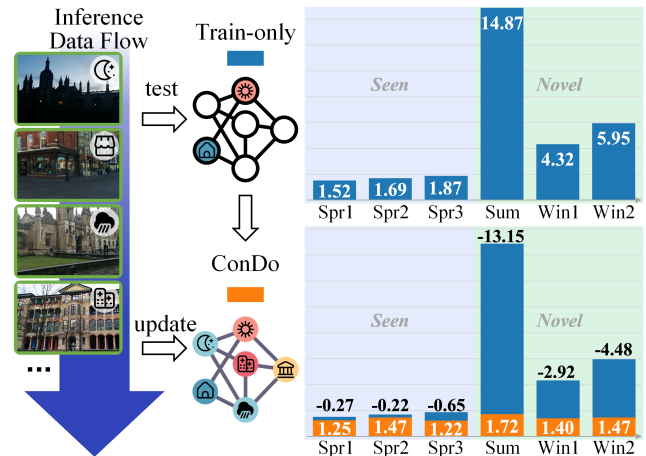


Figure 1: Teaser. We propose Continual Domain Expansion (ConDo) for APR, which utilizes unlabeled data seen during inference to expand the generalization domain of APR. Novel benchmarks are proposed to study practical scenarios where images are captured at novel poses or continually changing environments (left). The x-axis of histograms represents test data from various scans and y-axis indicates the estimated position median error. Trained only on data from spring, the deployed APR cannot handle summer and winter data (top). ConDo updates the model continually with unlabeled inference data and limited computation budgets, effectively expanding the generalization domain over time (bottom).

applying explicit geometric optimizations to obtain more fine-grained poses (Sarlin et al. 2019; Hyeon, Kim, and Doh 2021; Kim, Koo, and Kim 2023). Though more accurate, geometric optimization introduces computational overheads, making them limited when facing real-time applications. Absolute pose regression (APR) (Kendall, Grimes, and Cipolla 2015; Brahmhatt et al. 2018) is an important type of visual localization methods. It trains a light-weight scene-dependent neural network to directly output the camera pose of the query image. Such direct image-to-pose mapping makes APR highly efficient for both computation and memory (see Appendix A.1 for comparisons), suitable for real-time applications on edge devices. Comparing to multi-view methods like SLAM (Campos et al. 2021), APR can derive camera poses from a single

image. Though with clear advantages, the scene-dependent model training also limits the robustness of APR on novel data seen during inference (Sattler et al. 2019). The novel data can be captured either at poses distant from the training data (Sattler et al. 2019), or with unseen lighting, weather, or geometry (new construction) conditions caused by the change of time (Cai and Müller 2023). Fig. 1 shows an example (top) where the model only trained on data captured in spring sees inference data from summer and winter. The accuracy dropped heavily even though different trajectories have similar pose distributions.

A naive solution for this problem is to obtain new data with ground truth (GT) that cover novel poses and scene conditions, train a new model from scratch on both historical and new data, and then deploy the new model for inference. However, obtaining ground-truth data for APR often requires manual scene traverses with 3D scanners, which not only introduces extra laboring costs but also cannot guarantee to cover all novel data in a continually changing environment. Meanwhile, re-training models with more data needs more computation and time to converge.

In this work, we propose *Continual Domain Expansion (ConDo)* for APR. ConDo leverages unlabeled data seen after model deployment to continually and efficiently update APR. Though unsupervised domain adaptation methods have been proposed for standard classification/regression tasks, as shown later in the experiments, they struggle to generate effective supervision signals for APR. Inspired by the fact that scene-agnostic methods (Arandjelovic et al. 2016; Sartin et al. 2019; Von Stumberg and Cremers 2022; Campos et al. 2021) are much more robust to scene and pose changes, we instead generate supervision signals on unlabeled data by distilling knowledge from them. As shown in Fig.1, this simple yet effective strategy improves not only the performance on data from the same domain, but also the general robustness of APR, leading to better performance on other domains. Meanwhile, the model is updated continually without re-training, so that the computation does not grow over time. Besides the case of a single scene, the multi-head architecture is applied to make ConDo also applicable to sequentially revealed new scenes with a minimal model parameter increase. To thoroughly evaluate APR on data with practical changes, we construct benchmarks that cover 1) indoor and outdoor scenes, 2) large-scale city-level data, 3) (long-term) scene changes and novel camera poses.

Experiments validate the effectiveness of ConDo on different baseline architectures and data with both scene and pose changes. It reduces the localization error by more than an order of magnitude on challenging data. Comprehensive analysis shows the robustness of ConDo w.r.t. the knowledge distillation teacher, replay buffer sizes, compute budgets and so on. Comparing to model re-training, ConDo can reach similar performance with up to 25x compute/time reduction.

2 Related Work

Absolute pose regression. APR is a classical visual localization approach, which directly regresses the camera’s pose based on a single input image when revisiting a known environment. (Kendall, Grimes, and Cipolla 2015) proposed

the first APR method, which contains a feature extractor and pose regressor in the architecture. Follow-up methods improve the performance by introducing attention layers (Wang et al. 2020), Transformers (Shavit, Ferens, and Keller 2021) and Diffusion models (Wang et al. 2023). To better leverage scene information, visual odometry and motion constraints (Brahmbhatt et al. 2018; Xue et al. 2019) have been introduced. Recently, NeRFs (Neural Radiance Fields) have been used to generate more data (Moreau et al. 2022; Chen et al. 2022) or geometric constraints (Chen, Wang, and Prisacariu 2021; Moreau et al. 2023) for APR training. Though efficient, APR struggles to generalize to novel poses (Sattler et al. 2019) and scene changes (Cai and Müller 2023). ConDo is designed to address this problem.

Continual learning and other related problems. Conventional continual learning methods (Kirkpatrick et al. 2017; Aljundi, Chakravarty, and Tuytelaars 2017) aim to prevent catastrophic forgetting with limited storage. Recent approaches (Cai, Sener, and Koltun 2021; Prabhu et al. 2023) switch the focus to limited computation, aiming to achieve fast adaptation under practical limitations of training resources. This setup is similar to ConDo except that the ground truth labels are assumed to be available during continual model updates, which is impractical for localization systems that require high-end scanning devices to obtain accurate labels. Unsupervised domain adaptation (UDA) (Chen et al. 2021; Nejjar, Wang, and Fink 2023) aims to adapt a pre-trained model to a target domain without ground truth. Forgetting and computation budgets are not the major concern. Meta-learning (Finn, Abbeel, and Levine 2017) trains on diverse tasks to adapt with GT labels during inference, both of which are difficult to obtain for APRs. ConDo aims to continually adapt to new domains while preserving the performance of old ones. ConDo distills knowledge from scene-independent localization methods, which is more effective than standard UDA methods for APR.

3 Method

3.1 Preliminaries

Given an image $\mathbf{I} \in \mathbb{R}^{H \times W \times C}$, APR (Kendall, Grimes, and Cipolla 2015) learns a function $\mathbf{t}, \mathbf{r} = f(\mathbf{I}|\theta)$ parametrized by the neural network weights θ that maps \mathbf{I} to the camera position $\mathbf{t} \in \mathbb{R}^3$ and orientation $\mathbf{r} \in \mathbb{R}^4$ in a pre-defined scene Ω . Ω is defined by a set of training images $\mathcal{S}^\Omega = \{\mathbf{I}_i^\Omega\}_{i=1}^N$ with known poses $\mathcal{P}^\Omega = \{\mathbf{t}_i^\Omega, \mathbf{r}_i^\Omega\}_{i=1}^N$. The function f is a neural network commonly comprised of a feature extractor g and a regressor h , i.e., $f = h \circ g$ where g extracts the image level feature and h projects the extracted feature to \mathbf{t} and \mathbf{r} . Conventional APR frameworks train models on \mathcal{S}^Ω and \mathcal{P}^Ω with the regression loss (Kendall and Cipolla 2017):

$$\mathcal{L}(\mathbf{I}, \mathbf{t}^*, \mathbf{r}^*) = \|\mathbf{t} - \mathbf{t}^*\| e^{-s_t} + s_t + \|\mathbf{r} - \frac{\mathbf{r}^*}{\|\mathbf{r}^*\|}\| e^{-s_r} + s_r, \quad (1)$$

where \mathbf{t} and \mathbf{r} are the predicted pose on \mathbf{I} , $(\mathbf{t}^*, \mathbf{r}^*) \in \mathcal{P}^\Omega$ are the ground truth, and s_t and s_r are learnable parameters to balance the position and orientation losses. After training, the APR model is deployed to the environment for inference.

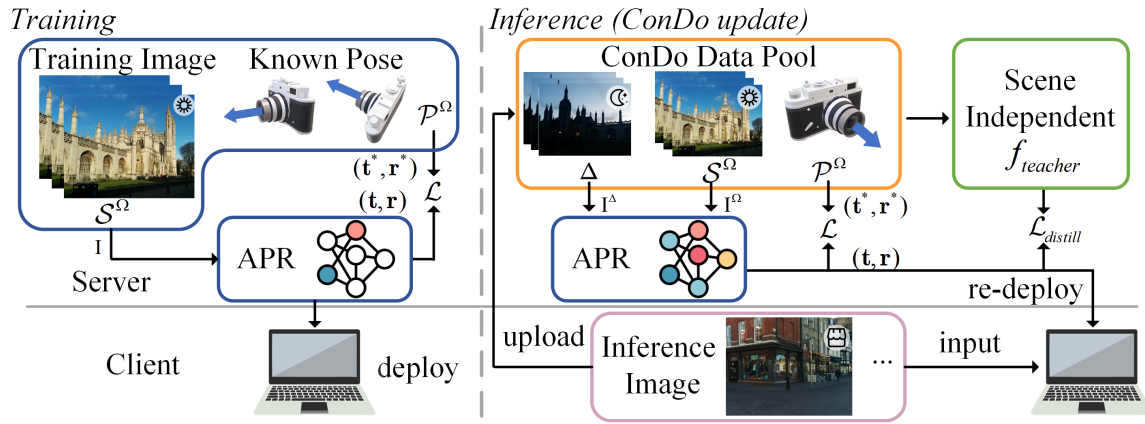


Figure 2: ConDo Pipeline. Left: After the normal APR training on labeled data, the model is deployed to the client. Right: After deployment, the client uploads the unlabeled data to the server. The server continually expands the generalization domain of APR by updating it with the labeled training data (\mathcal{S}^Ω , \mathcal{P}^Ω), unlabeled data Δ and a scene-independent teacher method f_{teacher} for knowledge distillation. Limited computation is assigned to each round of model update to ensure practical efficiency.

3.2 Continual Domain Expansion (ConDo)

Due to the scene-dependent nature, the deployed APR model cannot generalize well to images that have highly different poses or scene conditions compared to the training data \mathcal{S}^Ω . The key idea of Continual Domain Expansion (ConDo) is to continually update the APR model using the *unlabeled* data seen naturally after model deployment, so that the model can generalize to more novel poses and scene conditions over time by simply running in the environment.

As shown in Fig. 2, ConDo starts from the model θ trained on $(\mathcal{S}^\Omega, \mathcal{P}^\Omega)$. After model deployment, the (potentially multiple) clients, which use the newest model θ to perform localization, upload (asynchronously) the observed images to the server. The server collects all newly received images at time step k . These images are added to the pool of unlabeled data $\Delta = \{\mathbf{I}_j^\Delta\}_{j=1}^M$ and the current model θ_k is updated asynchronously from the clients using $\mathcal{S}^\Omega \cup \Delta$, with *limited computation* that is much less than model re-training. θ_k is re-deployed to the client after the update is finished.

In the main experiment of Sec. 5, we impose the constraint so that the compute for the pre-extended model training plus all update rounds of ConDo is the same as training one APR model from scratch on $\mathcal{S}^\Omega \cup \Delta$. This ensures that each ConDo update round uses much less compute and time compared to model re-training, so that it can be applied to life-long scenarios. We also experiment with various fixed computation budgets to validate the effectiveness of ConDo in applications with different resource limits.

To expand the generalization domain of APR without forgetting, we *uniformly* sample images from $\mathcal{S}^\Omega \cup \Delta$ to form a training batch during the model update. Though unsupervised domain adaptation (UDA) methods (Chen et al. 2021; Nejjar, Wang, and Fink 2023) have been proposed for standard image classification and regression, empirically (Sec. 5.2) they are not effective for APR. To generate effective supervision on unlabeled data Δ , we opt for a distillation-based approach. Inspired by the fact that scene-independent methods (Arand-

jelovic et al. 2016; Sarlin et al. 2019; Von Stumberg and Cremers 2022), though slower and more memory consuming during inference, are much more robust than APR to novel poses and scene conditions, we distill the knowledge from these methods to APR using Δ , so that the inference model can still maintain the memory and computation efficiency. Specifically, given a scene-independent method $f_{\text{teacher}}(\cdot)$, and a batch of data $\mathcal{B}^\Omega \cup \mathcal{B}^\Delta$ sampled from $\mathcal{S}^\Omega \cup \Delta$, the training objective is

$$\underset{\theta}{\text{minimize}} \frac{\sum_{\mathbf{I}^\Omega \in \mathcal{B}^\Omega} \mathcal{L}(\mathbf{I}^\Omega, \mathbf{t}_{\mathbf{I}^\Omega}^*, \mathbf{r}_{\mathbf{I}^\Omega}^*) + \sum_{\mathbf{I}^\Delta \in \mathcal{B}^\Delta} \mathcal{L}_{\text{distill}}(\mathbf{I}^\Delta, f_{\text{teacher}})}{|\mathcal{B}^\Omega| + |\mathcal{B}^\Delta|}, \quad (2)$$

where $\mathbf{t}_{\mathbf{I}^\Omega}^*, \mathbf{r}_{\mathbf{I}^\Omega}^*$ are the ground truth pose of \mathbf{I}^Ω . We choose HLoc (Sarlin et al. 2019) as the default f_{teacher} , with scene map built on $(\mathcal{S}^\Omega, \mathcal{P}^\Omega)$ (See Table.5 for the robustness of ConDo with other teachers). We set $\mathcal{L}_{\text{distill}} = \mathcal{L}(\mathbf{I}^\Delta, f_{\text{teacher}}(\mathbf{I}^\Delta))$, i.e., substituting the output of f_{teacher} into Eq. (1). As shown later in Sec. 5, this simple yet effective loss is sufficient to approach the performance of training with ground truth poses on Δ , and is robust to the choice of f_{teacher} . Meanwhile, distilling knowledge on data from new domains not only benefits the performance on the same domain, but can also improve the general robustness of APR, especially under scene condition changes.

By default, we assume the server has sufficient storage to maintain all historical data. For applications with limited server storage, we apply *reservoir sampling* (Rebuffi et al. 2017) to update the replay buffer. Given a sequence of K images and a storage \mathcal{M} sufficient to maintain N images, we push the first N images to the storage as usual. For the i -th image where $i > N$, we generate a random integer α between $[1, i]$. If $\alpha \leq N$, we replace the α -th image stored in \mathcal{M} with the i -th image in the sequence. Otherwise, we drop the i -th image. As shown later in Sec. 5.2, ConDo with reservoir sampling is robust to the replay buffer size.

For architectures capable of handling multiple scenes (Shavit, Ferens, and Keller 2021, 2023), ConDo

can also be applied when training data of new scenes and the inference data of old scenes arrive sequentially and interchangeably. When training data of a new scene is available during ConDo, we simply add an extra regression head to cope with new scene coordinates and perform normal ConDo training following Eq. (2). The only difference is that now the replay data of ConDo are sampled from all observed scenes. This strategy ensures that APR can handle both data from the same scene and the sequentially revealed multiple scenes, with a minor model parameter increase.

4 Benchmark

To thoroughly evaluate ConDo in practical scenarios, we construct large-scale benchmarks covering both the change of *scene conditions* (lighting, weather, season) and *camera poses*. Specifically, we collect public datasets with multiple rounds of scans of the same scene. To simulate the practical scenario, we split multiple scans of the same scene into training and inference and reveal the inference scans sequentially, i.e., every round of ConDo model update starts when a new inference scan is revealed. We randomly hold out $\frac{1}{8}$ images in each scan (training and inference) and use them to evaluate the generalization of APR on the corresponding scan. To create challenging evaluation data, instead of holding out individual images uniformly distributed in each scan, we hold several sets of images where each set is a continuous trajectory of the scan consisting of 16 images (see Fig. 3). The held-out evaluation data allow us to fully evaluate APR on images unseen both during normal training and ConDo.

To simulate *novel poses and the sequentially revealed multiple scenes*, we adopt standard APR datasets, namely, *7Scenes* and *Cambridge* (Glocker et al. 2013; Kendall, Grimes, and Cipolla 2015). These two datasets represent the case of indoor and outdoor scenes respectively and different scans of the same scene contain distinct trajectories, which are suitable to evaluate the case of novel poses. We adopt the same training and inference split as in the baseline APR methods (Kendall, Grimes, and Cipolla 2015; Shavit, Ferens, and Keller 2021). Please refer to Appendix A.2 for detailed information about the train/inference scan split, multi-scene revealing order, the used coordinate system, etc.

The drawbacks of *7Scenes* and *Cambridge* are the limited scene scale ($< 140m \times 40m$) and scene condition change. The lighting and weather conditions of both datasets remain similar in different scans, and there are no obvious *long-term scene changes (seasonal)* observed. To address these issues, we utilize large-scale driving datasets with both significant lighting changes (daytime to night time) and long-term scene changes (spring to winter). Specifically, we take the *Office Loop* and *Neighborhood*, which are two large-scale scenes in *4Seasons* (Wenzel et al. 2021) with a sufficient amount of scans (> 6) in the same scene. Each scan in these two scenes has a $> 2km$ trajectory spanned at multiple city blocks, which is much larger than conventional APR datasets. See Fig. 4 for sample images from different scene scans and Appendix A.2 for the concrete train/inference scan split.

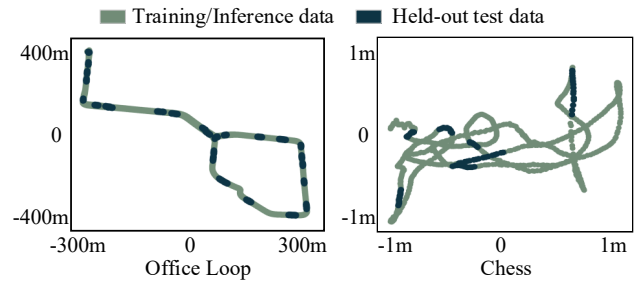


Figure 3: Data split visualization. $\frac{1}{8}$ images in each scan (training and inference) are held out for evaluations. To create challenging evaluation data, We randomly hold several sets of images where each set is a continuous trajectory of the scan consisting of 16 images. Left: Outdoor *Office Loop* data. Right: Indoor *Chess* scene in *7Scenes*.




Figure 4: *Office Loop* images. Obvious differences exist between training (Spring Sunny) and inference scans, e.g., overexposure (Summer Sunny), snow (Winter Snowy) and moving objects (Winter Sunny).

5 Experiments

Architectures. We validate ConDo on two representative APR architectures, namely PoseNet (PN) (Kendall, Grimes, and Cipolla 2015) and Pose-Transformer (PT) (Shavit, Ferens, and Keller 2021, 2023), which covers respectively the classic APR architectures for single and multiple scenes.

Implementation. Unless otherwise stated, the code and hyper-parameter settings of the baselines strictly follow the official code release. The original Pose-Transformer can use multiple regression heads and scene-dependent latent embeddings to handle multiple scenes. We only apply multiple regression heads since it is sufficient to achieve similar performance (Appendix A.4). APRs are first learned on training data until converging in the initial training. In the main experiment, we follow the setup of large scale continual learning (Cai, Sener, and Koltun 2021) and limit the computation budget of ConDo by first identifying the budget $b = \text{epoch} * \text{iteration_per_epoch} * \text{batch_size} / |\mathcal{S}^\Omega|$ for the baseline APR model to converge on the initial training data \mathcal{S}^Ω . b represents the average number of iterations required per image. Then for every round of ConDo update with N images newly revealed, we assign $N * b / \text{batch_size}$ training iterations (see Appendix A.2 for actual numbers of b) with the same batch size as the initial training, so that the whole ConDo procedure including initial training and all ConDo updates, consumes roughly only the budget to train one APR model from scratch on all revealed data. This ensures that we use *much less* computation than model re-training in every

Model	Strategy	Office Loop				Neighborhood			
		Training scan held-out data		Inference scan held-out data		Training scan held-out data		Inference scan held-out data	
		Median	Mean	Median	Mean	Median	Mean	Median	Mean
PN	1.Train-only	2.03/0.60	2.54/1.08	18.10/2.68	100.42/16.42	1.19/0.79	1.57/1.33	10.99/3.74	27.67/13.88
	2.ConDo	1.66/0.25	2.03/0.37	2.16/0.52	2.61/0.88	1.12/0.33	1.36/0.68	1.14/0.45	1.39/0.59
	3.Re-train with GT	1.64/0.20	2.05/0.30	1.80/0.19	2.42/0.51	0.99/0.26	1.23/0.37	1.00/0.23	1.29/0.31
	Improvement (1-2)	0.37/0.35	0.51/0.71	15.94/2.16	97.81/15.54	0.07/0.46	0.21/0.65	9.85/3.29	26.28/13.29
PT	1.Train-only	1.70/0.29	1.82/0.84	6.12/16.14	42.15/43.24	1.22/0.35	1.33/0.75	2.99/1.33	17.69/23.64
	2.ConDo	1.34/0.21	1.45/0.49	1.50/0.49	1.86/1.31	0.87/0.24	0.94/0.40	0.89/0.38	1.04/0.50
	3.Re-train with GT	1.41/0.19	1.39/0.64	1.46/0.18	1.59/0.63	0.73/0.22	0.77/0.36	0.76/0.19	0.84/0.33
	Improvement (1-2)	0.36/0.08	0.37/0.35	4.62/15.65	40.29/41.93	0.35/0.11	0.39/0.35	2.10/0.95	17.65/23.14

Table 1: Results on scene condition changes. Position (m) / orientation ($^\circ$) errors are reported for various strategies and architectures. Results on training and inference scans are reported separately for better analysis. ConDo improved the deployed APR (*Train-only*) by a large margin across architectures and scenes, reaching near upper bound *Re-train with GT* performance with limited computation budgets for model updates.

Model	Strategy	7Scenes				Cambridge			
		Training scan held-out data		Inference scan held-out data		Training scan held-out data		Inference scan held-out data	
		Median	Mean	Median	Mean	Median	Mean	Median	Mean
PN	1.Train-only	0.023/0.925	0.027/1.236	0.303/9.536	0.374/15.228	0.933/2.722	1.284/4.889	1.405/3.292	2.280/4.778
	2.ConDo	0.069/2.198	0.078/2.460	0.080/2.641	0.097/3.042	1.259/3.131	1.711/5.573	1.052/2.612	1.518/3.560
	3.Re-train with GT	0.024/0.995	0.029/1.265	0.024/1.016	0.028/1.191	0.786/2.446	1.228/4.896	0.864/2.155	1.288/2.797
	Improvement (1-2)	-0.046/-1.273	-0.051/-1.224	0.223/6.895	0.277/12.186	-0.326/-0.409	-0.427/-0.684	0.353/0.680	0.762/1.218
PT	1.Train-only	0.021/1.158	0.025/1.547	0.198/8.373	0.284/12.162	0.806/2.502	1.084/5.756	1.101/2.682	1.877/3.536
	2.ConDo	0.049/1.581	0.054/1.885	0.065/2.113	0.077/2.381	0.816/2.555	1.170/6.608	0.696/2.106	1.134/3.470
	3.Re-train with GT	0.026/1.267	0.030/1.535	0.026/1.261	0.030/1.507	0.709/2.228	1.041/4.738	0.690/2.012	1.030/2.550
	Improvement (1-2)	-0.028/-0.423	-0.029/-0.338	0.133/6.260	0.207/9.781	-0.010/-0.053	-0.086/-0.852	0.405/0.576	0.743/0.066

Table 2: Results on pose changes. Position (m) / orientation ($^\circ$) errors are reported for various strategies and architectures. Though seeing data with novel poses and from other scenes might not always benefit the performance on historical data, ConDo still consistently and significantly improved the generalization on inference scans.

round of ConDo update. See Sec. 5.2 for the comparison of ConDo and model re-training with varied computation budgets. All models are trained using one RTX-4090 GPU.

Evaluation protocol. Following the standard (Kendall, Grimes, and Cipolla 2015), we compute the median/mean camera position (m) and orientation ($^\circ$) error for different methods. For baselines, we train the model on the initial training data and evaluate the performance on all held-out test data (from both the training and inference scans). For ConDo, we first train the model over the initial training data, perform ConDo updates sequentially on all inference scans, and then evaluate the final model on the held-out data.

5.1 Main Results

Table 1 and 2 show the main results on benchmarks constructed in Sec. 4 with respectively the scene condition (*Office Loop* and *Neighbourhood*) and pose (*7Scenes* and *Cambridge*) changes. For each architecture (PN and PT), we show the results of 3 training frameworks:

1. *Train Only*: Normal APR training on the initial data S^Ω . Representing the practical base APR performance.
2. *ConDo*: The proposed ConDo strategy.
3. *Re-train with GT*: Train an APR model from scratch until convergence (*infinite computation budget at any time*) on both the training and inference data ($S^\Omega \cup \Delta$), with the *GT* label on Δ provided. This setup estimates the best performance that ConDo can achieve.

See Sec. 5.2 for further comparisons between ConDo and standard UDA methods.

ConDo significantly improved the performance across baseline architectures and datasets. This shows the capability of ConDo to adapt to both scene condition change (*Office Loop* and *Neighbourhood*) and data from novel poses (*7Scenes* and *Cambridge*), and sequentially learn to localize in multiple scenes (*7Scenes* and *Cambridge*). Before ConDo, the mean error of the baselines was much larger than the median error on the inference scans of large-scale datasets. E.g., PT had 42.15m mean error vs 6.12m median error on *Office Loop*. This indicates the existence of catastrophically failing predictions (see Fig. 5 for visualizations). After ConDo, not only mean and median errors were reduced significantly (by 23x and 4x respectively), but also the difference between them became small. This change shows the significantly improved generalization of ConDo.

The performance difference between *ConDo* and *Re-train with GT* was small, even though 1) ConDo only used unlabeled inference data, and 2) used limited compute for model updates on sequentially revealed data. For example, *Re-train with GT* on *Office Loop* with PT used $\sim 120h$ to reach the reported performance and performed much worse with less compute (see Sec. 5.2), while each ConDo update round only took $\sim 20h$, i.e., achieving similar accuracy with only $\frac{1}{6}$ of the compute. Note that this difference will further increase over time with more data collected. Interestingly, *ConDo* performed marginally better *Re-train with GT* in *Office Loop*, it was because HLoc we used is very accurate in this dataset, especially in terms of translation (see Table. 5 for details).

Another interesting observation is that whether new data helps the general robustness of an APR model depends on

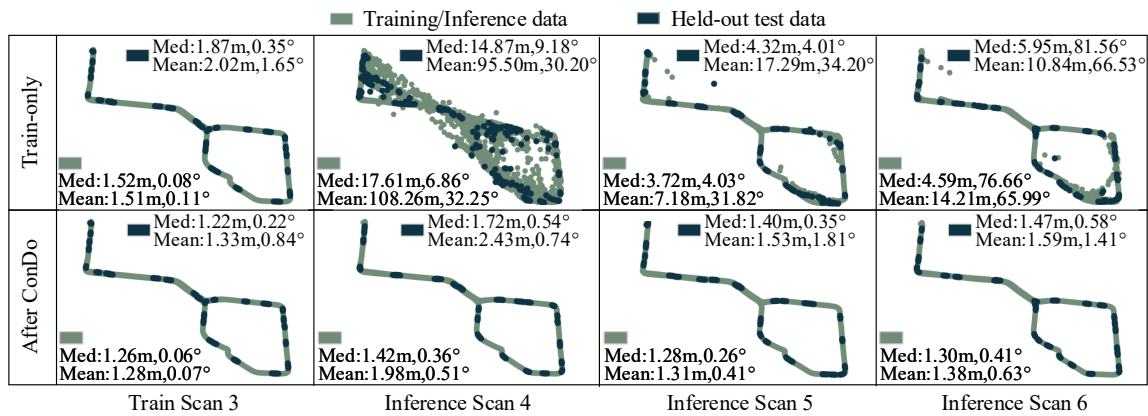


Figure 5: Result visualization on *Office Loop*. We visualize results on training and inference scans, where dark blue points indicate held-out test data and grey-green indicates training/inference data. Due to the space limit, we only visualize one training scan (*Train Scan 3*), see Appendix A.6 for other training scans. *Train-only* performed well on *Train Scan 3*, but cannot handle unseen scene condition changes (top row). By updating with unlabelled inference data, ConDo not only adapted to inference scans, but also generalized to the training ones (1.87m to 1.22m on the held-out data of *Train Scan 3*).

the type of change in the data. For data with scene condition changes (Table 1), learning from more data significantly improved the accuracy not only on the new scans, but also on previously seen scans. This result shows that training on more diverse data with different scene conditions helps to improve the general robustness of APR. On data with pose changes (Table 2), seeing new data may not always help the performance of old ones, even with ground truth labels (*Re-train with GT*). Appendix A.5 further analyzes the detailed reason for this phenomenon.

Fig. 5 visualizes the result of PT on *Office Loop*. Consistent with the conclusion from quantitative results, due to the weather/lighting condition change, the performance of APR dropped significantly on the inference scans, with many severe localization errors, especially in *Inference Scan 4*. After ConDo, these severe errors completely disappeared, and the localization accuracy not only improved on the inference scans, but also on the training ones (1.87m to 1.22m on the held-out data *Train Scan 3*). Fig. 6 shows the model accuracy on the held-out data of different scans after each round of ConDo update, which shows that seeing more data during ConDo can improve the *general robustness* of APR, resulting in a steady accuracy improvement. Note that the accuracy improvement on the training data was not caused by more training iterations in ConDo, since we have ensured that the model has converged during the *Train Only* phase, i.e., more training iterations without additional data would not help.

5.2 Analysis

This section analyzes the effectiveness of individual ConDo components, and we report results on *Office Loop* in the format of position (m)/orientation ($^\circ$) error. **ConDo vs UDA.** As mentioned in Sec. 3.2, unsupervised domain adaptation (UDA) is widely used to adapt models to novel data. In Table. 3, we compare ConDo with 3 most applicable UDA baselines, namely, *RSD* (Chen et al. 2021), *DARE* (Nejjar, Wang, and Fink 2023) and *MIC_{mae}* (Hoyer et al. 2023). Following

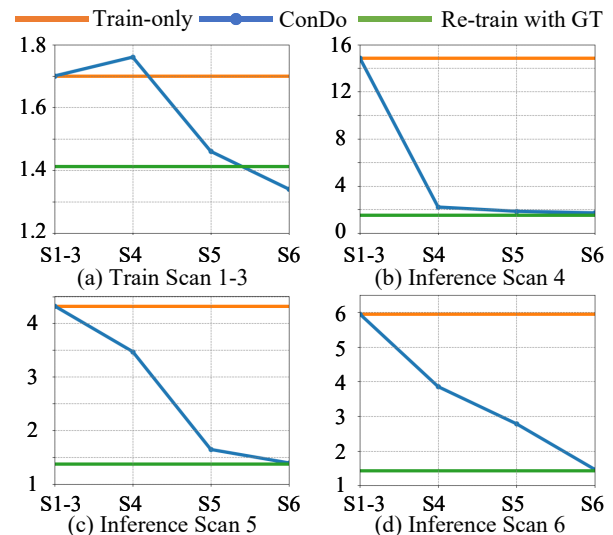


Figure 6: Intermediate ConDo performance. The median position error (m) of PT on *Office Loop* is reported. The x axis indicates the scans seen in each round of ConDo update. ConDo updates improved the accuracy not only on the current scan, but also on other scans. See Appendix A.7 for the version with *Train Scan 1-3* in separate figures.

the original papers, we reduce the feature dimension of RSD and DARE (from 1024) to 512 to avoid divergence. Empirically, RSD still diverges with this strategy, hence we report its result before divergence. We also run *Train-only*₅₁₂ with 512 dimension features to show the improvements of RSD and DARE. We also try *MIC_{moco}* which replaces masked inputs of *MIC_{mae}* with augmented ones. Since *MIC* is compatible with ConDo supervision, we also combine it with ConDo (*MIC_{mae/moco}*+ConDo). As shown in Table. 3, RSD and DARE have minor improvements over *Train-only*₅₁₂

Strategy	Train scan held-out data		Infer scan held-out data	
	Median	Mean	Median	Mean
Train-only ₅₁₂	2.00/0.29	2.20/0.74	6.79/12.77	55.28/44.41
RSD	1.99/0.30	2.09/0.80	6.65/12.52	53.55/43.30
DARE	1.86/0.29	1.98/0.73	6.72/13.08	54.19/44.13
Train-only	1.70/0.29	1.82/0.84	6.12/16.14	42.15/43.24
MIC _{mae}	2.59/0.73	3.03/2.19	6.29/39.87	20.03/50.99
MIC _{moco}	2.30/0.54	4.51/3.86	5.74/77.61	15.37/63.77
MIC _{mae} +ConDo	2.42/0.69	2.67/1.81	4.39/4.30	8.42/18.85
MIC _{moco} +ConDo	2.19/0.45	4.34/1.36	4.15/2.62	10.19/19.78
ConDo	1.34/0.21	1.45/0.49	1.50/0.49	1.86/1.31

Table 3: Comparison to UDA methods. We compare ConDo with RSD and DARE which are the latest UDA regression strategies, as well as the most effective and applicable strategy in UDA classification, MIC, and we modified its output from class distribution to regression space for the localization task.

Rate	Time	Strategy	Train scan held-out data		Infer scan held-out data	
			Median	Mean	Median	Mean
unlimited	120h	Re-train	1.41/0.19	1.39/0.64	1.46/0.18	1.59/0.63
1	20h	Re-train	1.86/0.24	1.90/0.65	1.79/0.44	2.01/1.19
		ConDo	1.34/0.21	1.45/0.49	1.50/0.49	1.86/1.31
1/2	10h	Re-train	2.03/0.36	2.07/0.84	2.12/0.57	2.30/1.47
		ConDo	1.56/0.24	1.72/0.61	1.64/0.50	1.92/1.39
1/4	5h	Re-train	2.75/0.61	2.97/1.20	2.70/0.65	3.11/1.58
		ConDo	1.81/0.27	1.88/0.64	1.91/0.52	2.34/1.49
1/8	2.5h	Re-train	3.46/0.78	3.64/1.56	3.42/0.87	3.68/1.98
		ConDo	1.95/0.35	2.03/0.69	2.23/0.62	2.51/1.64
1/100	12min	Re-train	8.13/2.28	9.38/7.36	9.12/2.62	11.49/8.03
		ConDo	2.41/0.47	2.42/2.53	2.89/0.91	3.31/2.80

Table 4: ConDo vs Re-train with varied training budgets. ConDo reached a similar accuracy up to 25x faster than Re-train.

and are far behind ConDo. MIC_{mae} and MIC_{moco} hurt both Train-only and ConDo.

Computation budget. Practical applications have varied computation budgets. To demonstrate the effectiveness of ConDo, we compare it with *Re-train with GT* (Re-train in short) under the same budget limits. Note that the *Re-train with GT* in the main results used unlimited compute, and took 120h for the last round of update, 6x higher than the ConDo update. As shown in Table.4, we gradually reduce the budget limit from 20h (the same as ConDo in the main results) to 12min ($\frac{1}{100}$ of the original budget). ConDo reached a similar accuracy much faster than *Re-train with GT*, even without using GT. E.g., the performance of ConDo with just 12min of model updates was on-par with *Re-train with GT* for 5h — a 25x compute/time reduction. Note that with only 20h, *Re-train* performed much worse than ConDo, even with GT. Appendix A.3 further shows the accumulated time of ConDo and *Re-train with GT* after all update rounds.

Replay buffer size. For applications with limited server storage, one can apply the reservoir buffer during ConDo. Following the standard approach (Cai and Müller 2023), we apply *reservoir sampling* to update ConDo and analyze its performance under different replay buffer sizes. As shown in Fig. 7, the performance ConDo only dropped slightly even under extreme replay buffer size limitation (10% of the overall dataset size), which still significantly improved over the *Train-only* baseline. This shows the effectiveness of ConDo

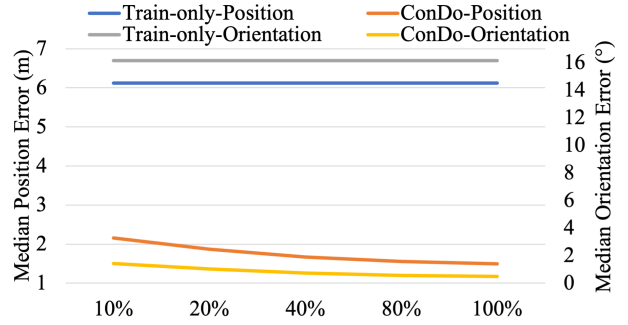


Figure 7: Effect of replay buffer sizes. The horizontal axis is the ratio between the replay buffer size and the whole dataset size. The vertical axis reports median errors in *Office Loop*.

$f_{teacher}(\cdot)$	Train held-out data		Infer held-out data		Teacher infer scan	
	Median	Mean	Median	Mean	Median	Mean
DM-VIO	1.50/0.29	1.67/0.55	5.21/0.81	5.56/1.17	5.08/0.82	5.74/0.84
ORB-3	1.40/0.20	1.43/0.42	3.39/0.94	3.63/1.24	3.15/0.93	3.14/0.97
NetVLAD	1.88/0.33	2.00/0.83	3.00/1.02	49.4/10.2	0.97/1.13	44.5/10.6
HLoc	1.34/0.21	1.45/0.49	1.50/0.49	1.86/1.31	0.05/0.31	0.41/0.51
GT	1.46/0.20	1.47/0.43	1.58/0.19	1.66/0.69	-	-

Table 5: Effect of teacher models. Results are evaluated on Office Loop and reported in the form of position (m) / orientation (°) errors. Left and middle column: results on held-out evaluation data of ConDos supervised by different teachers. Right column: teacher performance on the *whole set of* inference scans. Though slightly worse than HLoc, all other teachers can be effectively applied to ConDo and provide a reasonable performance improvement.

in practical applications with strong storage limits.

Teacher. Table.5 shows the performance of ConDo with different teachers $f_{teacher}(\cdot)$ (see Appendix A.8 for implementation details). The right-most column also reports the prediction error of each teacher model. The accuracy of ConDo and the teacher is positively correlated. On the other hand, despite the higher prediction noise, ConDo with other teachers all provided reasonable improvements to the base APR model.

Pre-training. Stronger pre-trained backbones often improve model generalization (Keetha et al. 2023; Käppeler et al. 2023). Appendix A.9 demonstrates that replacing the original APR backbones with stronger ones (Oquab et al. 2023) cannot replace the functionality of ConDo.

6 Conclusion

We have identified the problem of APR in generalizing to novel data during inference. We have proposed Continual Domain Expansion (ConDo) to address this problem. By distilling knowledge from scene-independent localization methods, ConDo allows APR to improve steadily and continually while running in deployed environments with unlabeled data. We have constructed large-scale benchmarks covering 1) indoor and outdoor scenes, and 2) the change of both environment conditions and novel poses. Experiments have verified the effectiveness and robustness of ConDo under varied teacher models, model architectures, scene types, compute budgets and replay buffer sizes.

References

- Aljundi, R.; Chakravarthy, P.; and Tuytelaars, T. 2017. Expert gate: Lifelong learning with a network of experts. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3366–3375.
- Arandjelovic, R.; Gronat, P.; Torii, A.; Pajdla, T.; and Sivic, J. 2016. NetVLAD: CNN architecture for weakly supervised place recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5297–5307.
- Brahmbhatt, S.; Gu, J.; Kim, K.; Hays, J.; and Kautz, J. 2018. Geometry-aware learning of maps for camera localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2616–2625.
- Cai, Z.; and Müller, M. 2023. CLNeRF: Continual Learning Meets NeRF. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 23185–23194.
- Cai, Z.; Sener, O.; and Koltun, V. 2021. Online continual learning with natural distribution shifts: An empirical study with visual data. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8281–8290.
- Campos, C.; Elvira, R.; Rodríguez, J. J. G.; Montiel, J. M.; and Tardós, J. D. 2021. Orb-slam3: An accurate open-source library for visual, visual–inertial, and multimap slam. *IEEE Transactions on Robotics*, 37(6): 1874–1890.
- Chen, S.; Li, X.; Wang, Z.; and Prisacariu, V. A. 2022. Dfnet: Enhance absolute pose regression with direct feature matching. In *European Conference on Computer Vision*, 1–17. Springer.
- Chen, S.; Wang, Z.; and Prisacariu, V. 2021. Direct-posenet: Absolute pose regression with photometric consistency. In *2021 International Conference on 3D Vision (3DV)*, 1175–1185. IEEE.
- Chen, X.; Wang, S.; Wang, J.; and Long, M. 2021. Representation Subspace Distance for Domain Adaptation Regression. In *ICML*, 1749–1759.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, 1126–1135. PMLR.
- Glocker, B.; Izadi, S.; Shotton, J.; and Criminisi, A. 2013. Real-time RGB-D camera relocalization. In *2013 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 173–179. IEEE.
- Hoyer, L.; Dai, D.; Wang, H.; and Van Gool, L. 2023. MIC: Masked image consistency for context-enhanced domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11721–11732.
- Hyeon, J.; Kim, J.; and Doh, N. 2021. Pose correction for highly accurate visual localization in large-scale indoor spaces. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15974–15983.
- Käppeler, M.; Petek, K.; Vödisch, N.; Burgard, W.; and Valada, A. 2023. Few-shot panoptic segmentation with foundation models. *arXiv preprint arXiv:2309.10726*.
- Keetha, N.; Mishra, A.; Karhade, J.; Jatavallabhula, K. M.; Scherer, S.; Krishna, M.; and Garg, S. 2023. Anyloc: Towards universal visual place recognition. *IEEE Robotics and Automation Letters*.
- Kendall, A.; and Cipolla, R. 2017. Geometric loss functions for camera pose regression with deep learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5974–5983.
- Kendall, A.; Grimes, M.; and Cipolla, R. 2015. PoseNet: A convolutional network for real-time 6-dof camera relocalization. In *Proceedings of the IEEE international conference on computer vision*, 2938–2946.
- Kim, M.; Koo, J.; and Kim, G. 2023. EP2P-Loc: End-to-End 3D Point to 2D Pixel Localization for Large-Scale Visual Localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 21527–21537.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13): 3521–3526.
- Moreau, A.; Piasco, N.; Bennehar, M.; Tsishkou, D.; Stanculescu, B.; and de La Fortelle, A. 2023. CROSSFIRE: Camera Relocalization On Self-Supervised Features from an Implicit Representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 252–262.
- Moreau, A.; Piasco, N.; Tsishkou, D.; Stanculescu, B.; and de La Fortelle, A. 2022. Lens: Localization enhanced by nerf synthesis. In *Conference on Robot Learning*, 1347–1356. PMLR.
- Nejjar, I.; Wang, Q.; and Fink, O. 2023. DARE-GRAM: Unsupervised domain adaptation regression by aligning inverse gram matrices. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11744–11754.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Prabhu, A.; Cai, Z.; Dokania, P.; Torr, P.; Koltun, V.; and Sener, O. 2023. Online continual learning without the storage constraint. *arXiv preprint arXiv:2305.09253*.
- Rebuffi, S.-A.; Kolesnikov, A.; Sperl, G.; and Lampert, C. H. 2017. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2001–2010.
- Sarlin, P.-E.; Cadena, C.; Siegwart, R.; and Dymczyk, M. 2019. From coarse to fine: Robust hierarchical localization at large scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12716–12725.
- Sattler, T.; Zhou, Q.; Pollefeys, M.; and Leal-Taixe, L. 2019. Understanding the limitations of cnn-based absolute camera pose regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3302–3312.

- Shavit, Y.; Ferens, R.; and Keller, Y. 2021. Learning multi-scene absolute pose regression with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2733–2742.
- Shavit, Y.; Ferens, R.; and Keller, Y. 2023. Coarse-to-Fine Multi-Scene Pose Regression with Transformers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Torii, A.; Arandjelovic, R.; Sivic, J.; Okutomi, M.; and Pajdla, T. 2015. 24/7 place recognition by view synthesis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1808–1817.
- Von Stumberg, L.; and Cremers, D. 2022. Dm-vio: Delayed marginalization visual-inertial odometry. *IEEE Robotics and Automation Letters*, 7(2): 1408–1415.
- Wang, B.; Chen, C.; Lu, C. X.; Zhao, P.; Trigoni, N.; and Markham, A. 2020. Atloc: Attention guided camera localization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 10393–10401.
- Wang, S.; Kang, Q.; She, R.; Tay, W. P.; Hartmannsgruber, A.; and Navarro, D. N. 2023. RobustLoc: Robust camera pose regression in challenging driving environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 6209–6216.
- Wenzel, P.; Wang, R.; Yang, N.; Cheng, Q.; Khan, Q.; von Stumberg, L.; Zeller, N.; and Cremers, D. 2021. 4Seasons: A cross-season dataset for multi-weather SLAM in autonomous driving. In *Pattern Recognition: 42nd DAGM German Conference, DAGM GCPR 2020, Tübingen, Germany, September 28–October 1, 2020, Proceedings 42*, 404–417. Springer.
- Xue, F.; Wang, X.; Yan, Z.; Wang, Q.; Wang, J.; and Zha, H. 2019. Local supports global: Deep camera relocalization with sequence enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2841–2850.