

HR-APR: APR-agnostic Framework with Uncertainty Estimation and Hierarchical Refinement for Camera Relocalisation

Changkun Liu¹, Shuai Chen³, Yukun Zhao¹, Huajian Huang¹, Victor Prisacariu³ and Tristan Braud^{1,2}

Abstract—Absolute Pose Regressors (APRs) directly estimate camera poses from monocular images, but their accuracy is unstable for different queries. Uncertainty-aware APRs provide uncertainty information on the estimated pose, alleviating the impact of these unreliable predictions. However, existing uncertainty modelling techniques are often coupled with a specific APR architecture, resulting in suboptimal performance compared to state-of-the-art (SOTA) APR methods. This work introduces a novel APR-agnostic framework, HR-APR, that formulates uncertainty estimation as cosine similarity estimation between the query and database features. It does not rely on or affect APR network architecture, which is flexible and computationally efficient. In addition, we take advantage of the uncertainty for pose refinement to enhance the performance of APR. The extensive experiments demonstrate the effectiveness of our framework, reducing 27.4% and 15.2% of computational overhead on the 7Scenes and Cambridge Landmarks datasets while maintaining the SOTA accuracy in single-image APRs.

I. INTRODUCTION

Camera relocalisation, estimating the six degrees of freedom (6-DoF) absolute camera pose in the world space, is a core component in many applications, including mobile robotics, navigation, and augmented reality. In recent years, Absolute Pose Regressors (APRs) have emerged as an appealing approach for monocular camera pose estimation. APRs use neural networks for directly inferring 6DoF camera poses from monocular frames and offer advantages in terms of computation and memory footprint over classical 3D structure-based methods [1], [2], [3], [4], [5]. However, APR methods often struggle with generalization beyond their training data, leading to inaccurate predictions [6].

Following the first APR model, PoseNet [7], numerous methods have been proposed to enhance the robustness and accuracy of pose predictions, including modifications to network architectures [8], [9], [10], [11], [12], augmentation of the training set with labelled synthetic images [13], [10], [6], different training strategies and loss functions [14], [15], [16]. Although these methods improve accuracy, they do not distinguish inaccurate predictions. In this study, we first re-examine prevailing APRs [7], [12], [16] and demonstrate that even such state-of-the-art (SOTA) APR methods still output

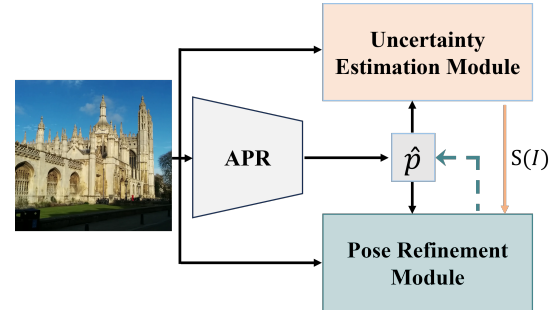


Fig. 1. HR-APR: APR-agnostic Framework with Uncertainty Estimation and Hierarchical Refinement.

significantly unreliable poses, raising the need for uncertainty estimation in real-life applications.

Uncertainty-aware (UA) APRs aim to distinguish unreliable predictions by providing additional uncertainty information with the estimated pose. Bayesian PoseNet [17] models the uncertainty by generating multiple hypotheses for each image at inference time. In AD-PoseNet [18], uncertainty is quantified via prior guided dropout. CoordiNet [19] models heteroscedastic uncertainty during training. Deng *et al.* [20], [21] and Zangeneh *et al.* [22] use pose distributions to represent uncertainty. However, these methods can be time-consuming [17], [18], and have weak extensibility as they rely on specific network architectures and specific training schemes [21], [22], [18], [19]. Furthermore, although these UA-APRs provide both pose predictions and uncertainty estimates, the accuracy of estimate poses is much lower than SOTA non-UA APRs [12], [16] that only output poses.

Such a gap between accuracy improvement and uncertainty estimation raises the need for more modular and more generic approaches to uncertainty estimation. In this paper, we propose HR-APR, a novel test-time APR-agnostic framework with uncertainty estimation and pose refinement (see Figure 1). The uncertainty estimation module integrates a new pose-based retrieval algorithm that fetches image feature embeddings in the training set. The cosine similarity between these retrieved features and the query image is then calculated to measure uncertainty. We leverage the uncertainty to optimize an iterative pose refinement pipeline [23].

We summarize our main contributions as follows:

- 1) We propose a novel APR-agnostic uncertainty estimation module to predict the uncertainty of the APR output during test time. This module integrates a new pose retrieval algorithm that fetches image feature embeddings in the training set. The cosine similarity

¹ Authors are with the Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, Hong Kong, {cliudg, yzhaoeg, hhuangbg}@connect.ust.hk, braudt@ust.hk

² Tristan Braud is also with the Division of Integrated Systems Design, The Hong Kong University of Science and Technology, Hong Kong

³ Shuai Chen and Victor Prisacariu are with the Active Vision Laboratory, Department of Engineering Science, University of Oxford, United Kingdom, {shuaic, victor}@robots.ox.ac.uk

between these retrieved features and the query image is then calculated to measure uncertainty.

- 2) We evaluate the accuracy of our uncertainty estimation module over three APR models with different architectures on indoor and outdoor datasets. The proposed method displays a clear correlation between pose error and uncertainty with similar performance across APR models, demonstrating the validity of the approach.
- 3) We further leverage the predicted uncertainty to reduce the overhead of an iterative pose refinement algorithm. HR-APR reduces the SOTA APR refinement pipeline’s overhead by 27.4% and 15.2% on the indoor and outdoor datasets, respectively, while maintaining the SOTA accuracy of single-image APR methods.

II. RELATED WORK

A. Absolute Pose Regression

APRs train neural networks for regressing the 6-DoF camera pose of query images. The seminal work in this area is introduced by PoseNet (PN) [7]. Further modifications to network architectures [8], [9], [10], [11], [12] and different training strategies [17], [14] have been made to improve accuracy and performance. MS-Transformer (MS-T) [12] follows MS-PN [24], which extends the single-scene APRs to multiple-scene APRs. Other approaches localise from image sequences [25], [26], [27], [28]. However, [6] has shown that most APR methods do not generalize well beyond training set via image retrieval baseline. To address this issue, [13], [10], [29] leverage additional synthetic training data. Other approaches [15], [16] adapt photometric or feature matching by applying unlabeled data with NeRF synthesis in a semi-supervised manner. However, using unlabeled data from the test set to finetune the APR network is impractical.

Most APR methods focus on improving pose prediction accuracy. However, we demonstrate in the following sections that not all pose predictions are reliable. The APRs mentioned above do not distinguish accurate predictions from poses with a large error. This paper introduces a modular, plug-and-play method to enable uncertainty estimation for non-UA APRs. Such uncertainty data can be used in the later stages of a visual positioning pipeline to improve tracking robustness and accuracy regardless of the underlying APR.

B. Uncertainty estimation

Several previous studies have investigated the pose predictions’ uncertainty in the training phase of APRs. In Bayesian PoseNet [17], uncertainty is captured by quantifying the variance among multiple inferences of the same input data using Monte Carlo Dropout. Similarly, AD-PoseNet [18] evaluates pose distribution by generating multiple hypotheses via prior guided dropout. CoordiNet [19] learns heteroscedastic uncertainty as an auxiliary task during the training. Poses and uncertainties output by CoordiNet are fused into an Extended Kalman Filter (EKF) to smooth the trajectories. Deng *et al.* [20], [21] and Zangeneh *et al.* [22] use pose distributions to represent uncertainty. While these UA APRs offer both pose predictions and uncertainty estimates, the accuracy of

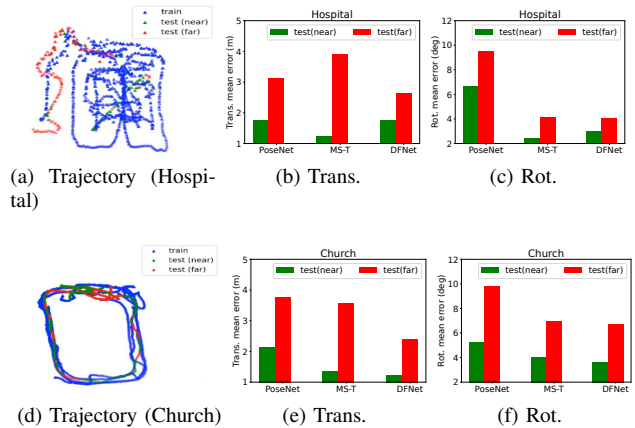


Fig. 2. The accuracy of APR predictions is highly corresponding to pose variants between query and training images. (a) and (d): Camera ground truth trajectory of Hospital and Church in Cambridge Landmarks dataset [7]. Blue: training set; Green: queries in the test set near the training set within 2 meters and 10 degrees; Red: queries in the test set far from the training set. (b), (c), (e), and (f) show that all three different APRs [7], [12], [16] have better predictions on test (near) than test (far).

estimate poses is much lower than SOTA non-UA APRs that only output poses, as shown in Section IV. Furthermore, existing UA APRs can be time-consuming [17], [18], require context-specific hyperparameters [21], or display weak extensibility as specific loss functions and modules need to be combined in a training scheme [21], [22], [18], [19].

Our framework enables greater flexibility regarding APR architecture and pose refinement compared to existing UA APRs. Mainstream pre-trained APRs can be integrated into our framework without modifying the network architecture or training schemes. The pose refinement module can leverage the uncertainty data to optimize computations significantly.

III. METHOD

Our objective is to enhance the robustness of APR methods and reduce the refinement overhead by identifying the reliabilities of predictions. Most existing APR methods suffer from limited generalizability due to overfitting their training sets. We first investigate several prevalent APR methods to show the accuracy of APR predictions is highly corresponding to pose variants between query and training images. The result is illustrated in Figure 2, where all three APR models exhibit more accurate predictions for test queries with viewpoints similar to those in the training set, compared to queries that fall outside the coverage of the training data. These quantitative results are also validated by the similar observation of [6], [30]. In light of this, we believe that the proximity of viewpoints also implies that these images have similar features. Therefore, we exploit pose predictions of APR to obtain the most corresponding features from the database. We then model the uncertainty of predictions by measuring the similarity between the features of the query image and the corresponding features of the database. The uncertainty information is further used as pose refinement constraints which can reduce the computational cost while achieving comparable performance.

A. Uncertainty estimation module

The uncertainty estimation module of HR-APR is shown in Figure 3. We build a database that stores low-dimensional feature embeddings of images from the training set and their associated ground truth poses, which operates as follows:

- 1) Given a query image I , APR P outputs estimated translation $\hat{\mathbf{x}}$ and rotation $\hat{\mathbf{q}}$ so that $P(I) = \hat{p} = \langle \hat{\mathbf{x}}, \hat{\mathbf{q}} \rangle$. All feature embeddings F_r with poses within ranges of threshold d_{th} to $\hat{\mathbf{x}}$ are retrieved. Then, proceed to step 2). If there is no valid pose in the database, we simply assume that the similarity is 0 and skip to step 3).
- 2) A feature extractor E extracts the feature embeddings of the query image I as f_q . Then, we compute the cosine similarity between f_q and each feature embedding $f_r^i \in F_r$, retrieved in step 1). We take the maximum value as the final similarity score.
- 3) The estimated pose that obtain a similarity score above threshold γ is deemed to be reliable, and thus to be output directly or refined with small steps. Otherwise, the poses are considered unreliable and to be treated with more refinement steps.

Compared to typical image retrieval approaches that often require large computational complexity in the visual feature descriptor matchings, our pose retrieval is a super cheap alternative by only relying on 3-D position searches with a significant reduction in the search space. The database consisting of pose-feature embedding pairs also saves a lot of memory footprint compared to the image database in [2].

B. CNN feature embeddings

Our feature extractor E is similar to an ordinary PoseNet [7], [14], which predicts a 6-DoF camera pose for an input image. E has an FC layer before the final regressor of feature size $n = 1024$. The difference is that E also outputs the FC layer as feature embedding. E is supervised by the learnable loss function [14]:

$$\mathcal{L}_\delta = \mathcal{L}_x \exp(-\hat{s}_x) + \hat{s}_x + \mathcal{L}_q \exp(-\hat{s}_q) + \hat{s}_q \quad (1)$$

, where $\mathcal{L}_x = \|\hat{\mathbf{x}} - \mathbf{x}\|_2$ and $\mathcal{L}_q = \|\frac{\hat{\mathbf{q}}}{\|\hat{\mathbf{q}}\|} - \mathbf{q}\|_2$. To measure the difference between two feature embeddings, we compute cosine similarity between feature embedding f_q extracted from query images and each retrieved feature embeddings $f_r^i, f_r^j \in F_r$.

$$\cos(f_q, f_r^i) = \frac{f_q \cdot f_r^i}{\|f_q\|_2 \cdot \|f_r^i\|_2} \quad (2)$$

, where $f_q, f_r^i \in \mathbf{R}^{1 \times 1024}$. The Similarity score of the query image I is:

$$S(I) = \max(\cos(f_q, f_r^i), f_r^i \in F_r) \quad (3)$$

The similarity score of the query image I is $S(I)$, where $-1 \leq S(I) \leq 1$, which reflects the reliability of the prediction.

C. Outcomes

The proposed uncertainty estimation module can identify an unreliable pose \hat{p} under two scenarios:

- 1) The training set does not contain an image close enough to the estimated location $\hat{\mathbf{x}}$.
- 2) Although feature embeddings of images in training set are retrieved from the database based on $\hat{\mathbf{x}}$, they all present too few similar features.

In the first scenario, \hat{p} can only get 0 similarity score because of the limited generalization ability of the neural network. The APR thus has a high chance of predicting a vastly incorrect pose for an image I far from the training set. In the second scenario, although valid most-similar feature embeddings F_r extracted from images in the training set is found based on $\hat{\mathbf{x}}$, all $f_r^i \in F_r$ are not similar enough to f_q . Either the orientations of the query image and these images are very different, leading to little overlap, or the predicted $\hat{\mathbf{x}}$ and $\hat{\mathbf{q}}$ have a large error, and F_r found according to this pose is a false positive.

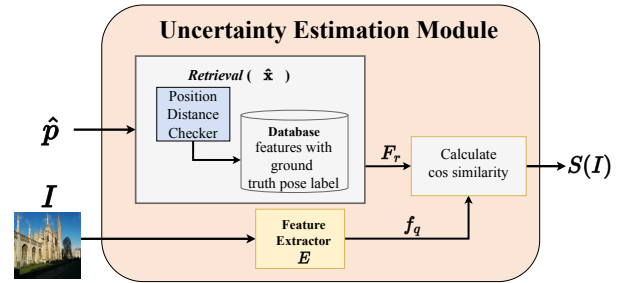


Fig. 3. Uncertainty estimation module.

D. Pose refinement module

In this paper, we incorporate the test-time Neural feature synthesizer (NeFeS) refinement pipeline [23] in our pose refinement module. We modify the refinement procedure to be uncertainty-aware. The refinement procedure is as follows: (i) For estimated camera pose \hat{p} with a similarity score in the uncertainty estimation module, NeFeS N renders a dense feature map m^{rend} given \hat{p} . (ii) At the same time, a feature map extractor G extracts a dense feature map $m^G = G(I)$ from the query image. (iii) The pose \hat{p} is iteratively refined by minimizing the feature cosine similarity loss [23] between m^{rend} and m^G . The refined steps depend on the similarity score $S(I)$ in our *uncertainty estimation module*.

E. Hyperparameters

During test time, the proposed method is controlled by two hyperparameters: the distance threshold d_{th} , and the similarity threshold γ . d_{th} depends on the size of the scene.

IV. EXPERIMENT

To demonstrate the scalability and effectiveness of HR-APR, we integrate the latest representative mainstream APR architectures into our framework and test them over indoor and outdoor datasets of various scales.

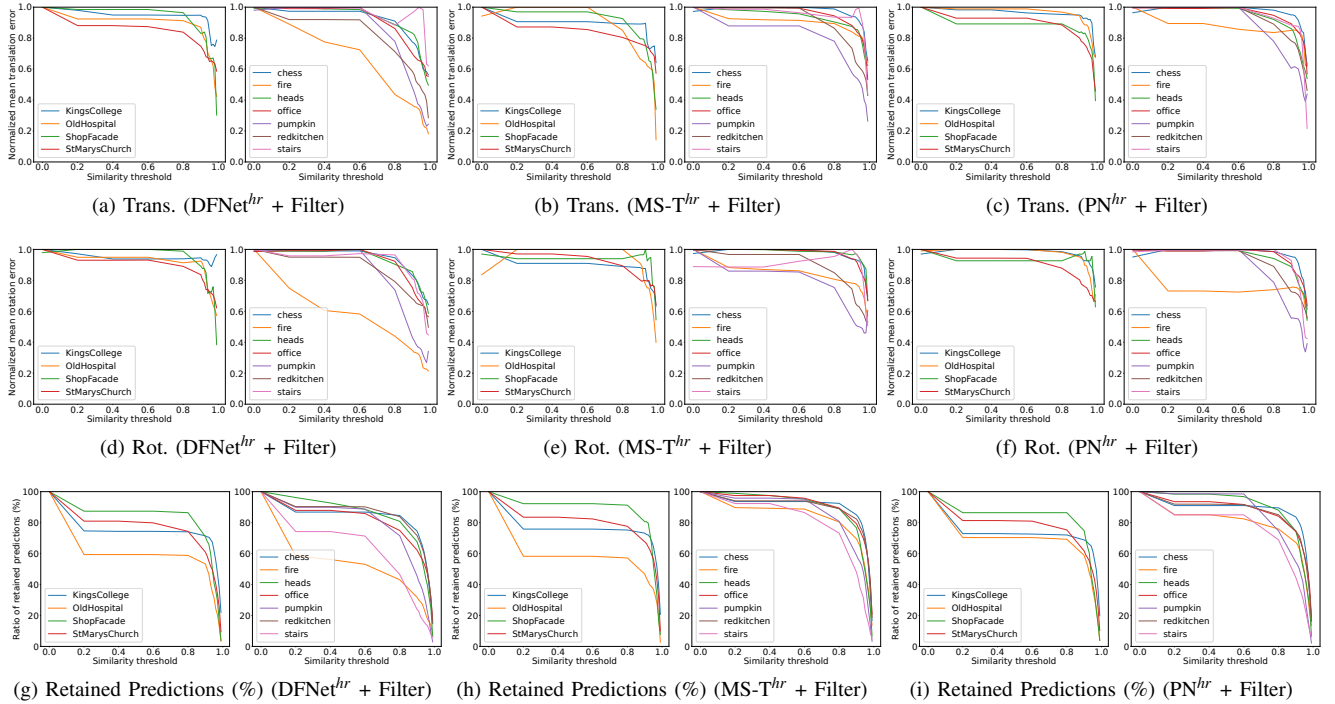


Fig. 4. Uncertainty evaluation on the 7Scenes and Cambridge Landmarks datasets. Subfigures (a)-(f) show the correlation between similarity threshold γ and normalized pose error (0-1). Based on the similarity threshold, uncertain samples with low similarity scores are gradually removed. Subfigures (g)-(i) in the last row show the ratio of retained predictions (%) as threshold increases. We observe that as we remove the samples with low similarity scores the overall error drops indicating a clear correlation between our predictions and the actual inaccurate predictions.

A. Datasets

The 7Scenes dataset [31], [32] is an indoor dataset consisting of seven small scenes from $1m^3$ to $18m^3$. Each scene contains a training set with 1000 to 7000 images and a test set with 1000 to 5000 images. We use the Structure from Motion (SfM) ground truth provided by [33] for our experiment. The Cambridge Landmarks [7] dataset represents six large-scale outdoor scenes ranging from $900m^2$ to $5500m^2$. We utilize four out of six scenes for comparative evaluation, and the SfM ground truth provided by [7]. Each scene contains 231 to 1487 images for training and 103 to 530 for testing.

B. Implementation Details

In this paper, we train the feature extractor E with EfficientNet-B0 [34] as the backbone in the *uncertainty estimation module*. The training process of E follows [7], [14]. It is trained using Adam optimizer [35] with an initial learning rate of $1e^{-4}$ and a weight decay of $5e^{-4}$. During training, all input images are resized to 256×256 and then randomly cropped to 224×224 .

We implement the *uncertainty estimation module* of HR-APR over three recent APR models: **PoseNet (PN)** [7] is the classic pose regression architecture. **MS-Transformer (MS-T)** [12] aggregates the activation maps with self-attention and queries the scene-specific information. **DFNet** [16] is trained by feature matching loss. We train these three models on our platform using the open-source codes. **We incorporate these APRs into our HR-APR framework, referring to them as APR^{hr} .** We set d_{th} as $0.2m$ for the 7Scenes dataset and $1.5m$

for the Cambridge dataset. We set γ range from $0.95 \sim 0.98$. Furthermore, we show the impact of γ in the next subsection.

In this paper, we implement the NeFeS refinement pipeline [23] in our pose refinement module. The term $APR^{hr} + NeFeS_{ls(y)}^{hs(x)}$ denotes the refinement process based on the similarity of estimated poses. We refine high similarity (hs) poses that are above the similarity threshold γ with x steps. Conversely, we refine low similarity (ls) poses that are below the similarity threshold γ with y steps. For instance, $APR^{hr} + NeFeS_{ls50}^{hs10}$ indicates that we refine poses with similarity scores higher than γ for 10 steps, but refine poses with a low similarity (ls) $< \gamma$ for 50 steps. Alternatively, we use $APR^{hr} + Filter$ to denote the scenario where we exclude the pose refinement module and directly filter out estimated poses with low similarity scores (ls). This approach helps us analyze the effectiveness of our framework, as shown in the uncertainty evaluation below. Besides, $APR^{hr} + Filter$ can be used to reject bad predictions in real mobile robotics applications simply.

C. Uncertainty Evaluation

We compare our method quantitatively with three mainstream single-frame non-UA APR methods and UA APR methods. We utilized the resulting similarity score to measure prediction uncertainty in this paper. Our *uncertainty estimation module* applies to almost all APRs, and we present results over three different models. We gradually increase the threshold γ , filter predictions whose similarity scores are below γ , and analyze the average error of the remaining

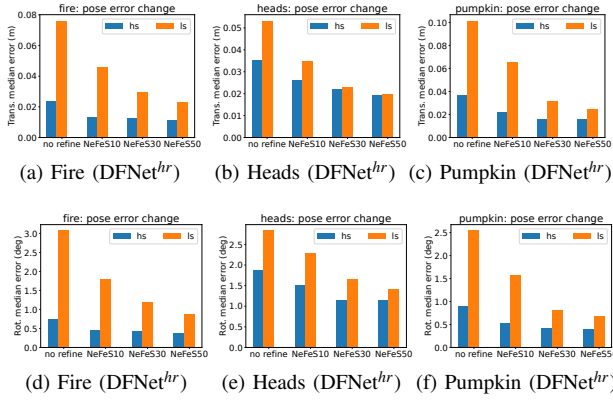


Fig. 5. Plots of translation and rotation errors against the number of iteration for images pass the similarity threshold (hs) and images with low similarity scores (ls) do not pass the similarity threshold on some scenes of 7Scenes. The hs predictions ratio of each scene is provided in Table I. NeFeS m denotes running the refinement process for m iterations.

TABLE I

PERCENTAGE OF PREDICTIONS WITH HIGH ($0.25M$, 2°), MEDIUM ($0.5M$, 5°), AND LOW ($5M$, 10°) ACCURACY [36] (HIGHER IS BETTER). THE VALUE IN PARENTHESES REPRESENTS THE RATIO OF RETAINED PREDICTIONS (HS) IDENTIFIED IN THE TEST SET BY OUR METHOD.

	Dataset	DFNet [16]	DFNet ^{hr} + Filter (ours)
7Scenes	Chess	81.1/97.3/100	87.6/99.6/100 (75%)
	Fire	46.5/75.7/94.5	92.3/97.1/100 (21%)
	Heads	44.8/87.2/98.1	56.6/98.4/100 (44%)
	Office	63.6/89.7/96.6	82.7/97.9/99.9 (31%)
	Pumpkin	51.6/77.6/93.1	80.5/99.1/100 (28%)
	Kitchen	59/82.1/94.5	77.9/99.6/100 (30%)
	Stairs	32.9/79.7/98.5	66/100/100 (14%)
Cambridge	Kings	4.1/27.1/96.5	5.3/28.9/100 (55%)
	Hospital	0/1/94.5	0/4.3/100 (13%)
	Shop	7.8/31.1/97.1	19/50/100 (16%)
	Church	0.9/9.1/87.9	1.9/14.6/99.1 (40%)

sample, as shown in Figure 4. The subfigures Figure 4.g-i indicate a decrease in the ratio of retained predictions as the threshold γ increases, coupled with the decrease in normalized mean translation and rotation errors of the retained predictions as γ increases in Figures 4.a-f. This clear correlation demonstrates that samples with higher similarity scores tend to yield more accurate estimations.

1) *7Scenes*: For all three APR models, when the similarity threshold γ is higher than 0.95, the translation and rotation mean errors are reduced by 25% to 80% across the scenes.

2) *Cambridge*: For all three APR models, when γ is higher than 0.95, the translation and rotation mean errors are reduced by 15% to 80% on different scenes.

Non-UA APRs (MS-T and DFNet) exhibit higher accuracy individually compared to UA APRs (Bayesian PN, AD-PN, BMDN, and VaPoR) as shown in Table II and Table III. Our new APR-agnostic *uncertainty estimation module* enables uncertainty estimation for such models, allowing us to leverage their greater accuracy in applications that require uncertainty measures.

D. UA pose refinement evaluation

Adding uncertainty awareness into APRs through the *uncertainty estimation module* enables filtering unreliable poses and prioritizing the optimization of less accurate predictions.

1) *7Scenes*: Table I illustrates the performance improvement provided by solely filtering unreliable poses identified by the pose estimation module. DFNet^{hr}+Filter achieves higher percentages in all accuracy levels by keeping the 14~75% predictions with reasonably high similarity scores pass γ . Furthermore, we completely filter out the pose predictions in the low accuracy level ($5m$, 10°) in six of seven scenes, and reduce it to 0.1% in the Office scene. Regarding pose optimization, Figure 5 shows that high similarity (hs) predictions, with similarity scores above γ , largely converge within 10 steps of optimization using the NeFeS method. On the other hand, low similarity (ls) predictions still require 10 to 50 steps to converge.

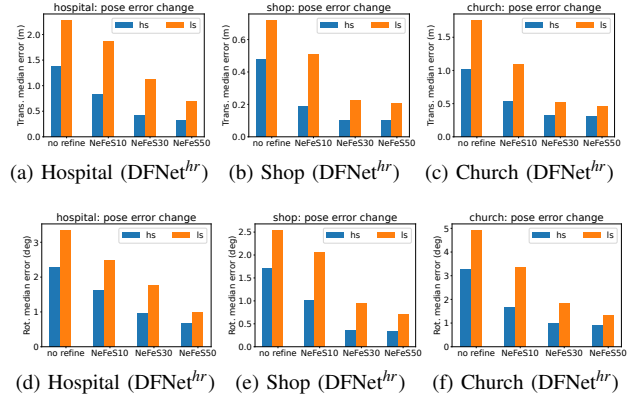


Fig. 6. Plots of translation and rotation errors against the number of iteration for images pass the similarity threshold (hs) and images with low similarity scores (ls) do not pass the similarity threshold on some scenes of Cambridge. The hs predictions ratio of each scene is provided in Table I. NeFeS m denotes running the refinement process for m iterations.

2) *Cambridge*: Similar to 7Scenes, Table I shows that filtering unreliable poses achieves higher percentages in all accuracy levels by keeping the 13~55% predictions with reasonably high similarity scores pass γ . Besides, we filter out the pose predictions in low accuracy level ($5m$, 10°) in three of four scenes and reduce it to 0.9% in the Church. As shown in Figure 6, high similarity (hs) predictions, with similarity scores above γ , largely converge within 30 steps of optimization using the NeFeS method, but low similarity (ls) predictions still require 30 to 50 steps to converge.

Table II and Table III show the accuracy of our UA pose refinement method compared to the SOTA non-UA APRs and refinement pipeline. The results demonstrate that our method achieves comparable accuracy to SOTA methods while reducing the optimization overhead by 27.4% and 15.2% on the indoor and outdoor datasets, respectively.

E. Analysis

Figure 2 shows that all three APR models have higher accuracy on query images closer to the training set's viewpoints. Without knowing the ground truth of the query image,

TABLE II

COMPARISONS ON 7-SCENES DATASET. THE MEDIAN TRANSLATION AND ROTATION ERRORS (M°) OF DIFFERENT METHODS AND THE AVERAGE REFINE STEPS FOR EACH QUERY IMAGE. THE BEST RESULTS ARE IN BOLD (LOWER IS BETTER).

	Methods	Chess	Fire	Heads	Office	Pumpkin	Redkitchen	Stairs	Avg. [m°]	Avg. Steps
UA APR	Bayesian PN [17]	0.37/7.24	0.43/13.7	0.31/12.0	0.48/8.04	0.61/7.08	0.58/7.54	0.48/13.1	0.47/9.81	-
	CoordiNet [19]	0.14/6.7	0.27/11.6	0.13/13.6	0.21/8.6	0.25/7.2	0.26/7.5	0.28/12.9	0.22/9.70	-
	BMDN ¹ [21]	0.10/6.47	0.26/14.8	0.13/13.4	0.19/9.73	0.20/9.40	0.19/10.9	0.34/14.1	0.20/11.26	-
	VaPoR [22]	0.17/6.9	0.30/14.1	0.17/14.5	0.24/9.3	0.30/8.3	0.26/10.2	0.47/15.5	0.27/11.26	-
non-UA APR	PN [7]	0.10/4.02	0.28/9.44	0.18/13.9	0.17/5.99	0.22/5.18	0.23/5.91	0.34/11.6	0.22/7.16	-
	MS-T [12]	0.11/6.45	0.23/10.9	0.13/13.1	0.18/8.18	0.16/6.87	0.17/8.44	0.30/10.4	0.18/9.19	-
	DFNet [16]	0.03/1.12	0.06/2.30	0.04/2.29	0.06/1.54	0.07/1.92	0.07/1.74	0.12/2.63	0.06/1.93	-
APR Refine	DFNet + NeFeS10 [23]	0.02/0.55	0.03/1.20	0.03/1.88	0.04/1.08	0.04/0.98	0.03/0.97	0.10/2.47	0.04/1.30	10
	DFNet + NeFeS50 [23]	0.02/0.57	0.02/0.74	0.02/1.28	0.02/0.56	0.02/0.55	0.02/0.57	0.05/1.28	0.02/0.79	50
Ours	DFNet ^{hr} + NeFeS ₅₀ ¹⁰	0.02/0.55	0.02/0.75	0.02/1.45	0.02/0.64	0.02/0.62	0.02/0.67	0.05/1.30	0.02/0.85	36.3

¹ Results of BMDN taken from [22].

TABLE III

COMPARISONS ON CAMBRIDGE DATASET. THE MEDIAN TRANSLATION AND ROTATION ERRORS (M°) OF DIFFERENT METHODS AND THE AVERAGE REFINE STEPS FOR EACH QUERY IMAGE. BEST RESULTS ARE IN BOLD (LOWER IS BETTER).

	Methods	Kings	Hospital	Shop	Church	Avg.	Avg. Steps
UA APR	Bayesian PN	1.74/ 4.06	2.57/ 5.14	1.25/ 7.54	2.11/ 8.38	1.92/6.28	-
	AD-PN	1.31/1.67	2.28/4.80	1.22/6.17	-/-	-/-	-
	BMDN ¹	1.51/2.14	2.25/3.93	3.52/5.41	2.16/5.99	2.36/4.37	-
	VaPoR	1.65/2.88	2.06/4.33	1.02/6.03	1.80/5.90	1.63/4.79	-
non-UA APR	PN	0.93/2.73	2.24/7.88	1.47/6.62	2.37/5.94	1.75/5.79	-
	MS-T	0.85/1.45	1.75/2.43	0.88/3.20	1.66/4.12	1.29/2.80	-
	DFNet	0.73/2.37	2.00/2.98	0.67/2.21	1.37/4.02	1.19/2.90	-
APR Refine	DFNet + NeFeS30	0.37/0.64	0.98/1.61	0.17/0.60	0.42/1.38	0.49/1.06	30
	DFNet + NeFeS50	0.37/ 0.54	0.52/0.88	0.15/0.53	0.37/1.14	0.35/0.77	50
Ours	DFNet ^{hr} + NeFeS ₃₀ ³⁰	0.36/0.58	0.53/0.89	0.13/0.51	0.38/1.16	0.35/0.78	42.4

¹ Results of BMDN taken from [22].

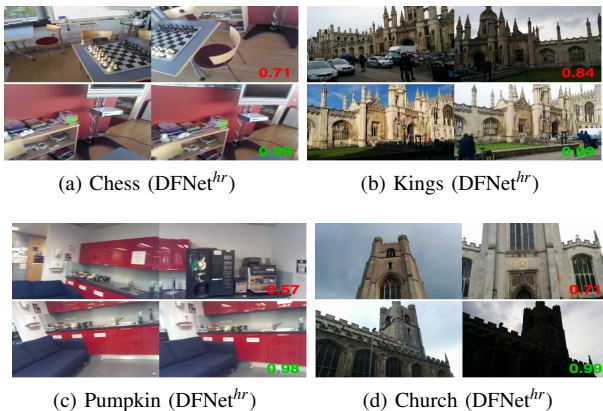


Fig. 7. Examples of retrieved image pairs in Cambridge and 7Scenes. For each pair, left side is the query image in test set, right side is the training set image retrieved by pose-based algorithm with the maximum similarity. The red numbers in the upper row of each subfigure represent the two images with low similarity and the green numbers in the lower row represent the two images with high similarity.

our uncertainty estimation module utilizes image features to measure the similarity between the query image and the training set, as we expect the similarity of image features to equate the proximity of viewpoints. We visualize the query image and the corresponding image in the training set retrieved by our pose-based algorithm with the maximum similarity. As shown in Figure 7, some query images are very similar to the training set images, while others are less

straightforward to match. The clear correlation in Figure 4 and improvement in all accuracy levels after filtering in Table I confirm our uncertainty estimation module is effective and fits the design outcomes in Section III-C. Besides, Figure 7.b and Figure 7.d show that the high-level and low-dimensional feature embeddings extracted from PoseNet-based feature extractor E are robust to different lighting, weather, moving objects, and pedestrians where point-based SIFT registration fails [7].

F. System Efficiency

We evaluate the processing time of the proposed framework on a PC equipped with an NVIDIA GeForce GTX 3090 GPU. Based on the implementation in Section III and Section IV, we repeat each measurement 1000 times. The feature extraction of E takes 3.4ms. Pose-based feature retrieval with $\mathcal{O}(n)$ complexity takes 1.98ms, while similarity calculation takes 1.3ms in the Heads dataset. Overall, the *uncertainty estimation module* only adds about an extra 6.7ms to each inference time of APRs. That means predictions can get the uncertainty in real-time. The feature extractor E with EfficientNetb0 as backbone only takes 20 MB to store the weights, and each 1×1024 feature embedding of an image only takes 4.2 KB. That means our *uncertainty estimation module* is also very storage efficient. Through its high modularity, our framework provides a new paradigm for real-life camera relocalisation with APRs. It can integrate most existing and future APRs and test-time refinement methods with minimal changes.

V. CONCLUSIONS

We introduce an APR-agnostic framework, HR-APR, which includes an uncertainty estimation module that uses a novel pose retrieval algorithm to calculate the cosine similarity between image feature embeddings in the training set and the query image. This uncertainty estimation is used to optimize an iterative pose refinement algorithm. We evaluate the performance of our method on indoor and outdoor datasets using three different APRs. Our results demonstrate a clear correlation between pose error and uncertainty, validating the effectiveness of our approach. Moreover, HR-APR significantly reduces the computational overhead of the refinement pipeline while maintaining accuracy.

REFERENCES

- [1] M. Dusmanu, I. Rocco, T. Pajdla, M. Pollefeys, J. Sivic, A. Torii, and T. Sattler, "D2-net: A trainable cnn for joint description and detection of local features," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 8092–8101.
- [2] P.-E. Sarlin, C. Cadena, R. Siegwart, and M. Dymczyk, "From coarse to fine: Robust hierarchical localization at large scale," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12716–12725.
- [3] H. Taira, M. Okutomi, T. Sattler, M. Cimpoi, M. Pollefeys, J. Sivic, T. Pajdla, and A. Torii, "Inloc: Indoor visual localization with dense matching and view synthesis," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7199–7209.
- [4] H. Noh, A. Araujo, J. Sim, T. Weyand, and B. Han, "Large-scale image retrieval with attentive deep local features," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 3456–3465.
- [5] T. Sattler, B. Leibe, and L. Kobbelt, "Efficient & effective prioritized matching for large-scale image-based localization," *IEEE transactions on pattern analysis and machine intelligence*, vol. 39, no. 9, pp. 1744–1756, 2016.
- [6] T. Sattler, Q. Zhou, M. Pollefeys, and L. Leal-Taixe, "Understanding the limitations of cnn-based absolute camera pose regression," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 3302–3312.
- [7] A. Kendall, M. Grimes, and R. Cipolla, "Posenet: A convolutional network for real-time 6-dof camera relocalization," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 2938–2946.
- [8] F. Walch, C. Hazirbas, L. Leal-Taixe, T. Sattler, S. Hilsenbeck, and D. Cremers, "Image-based localization using lstms for structured feature correlation," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 627–637.
- [9] I. Melekhov, J. Ylioinas, J. Kannala, and E. Rahtu, "Image-based localization using hourglass networks," in *Proceedings of the IEEE international conference on computer vision workshops*, 2017, pp. 879–886.
- [10] J. Wu, L. Ma, and X. Hu, "Delving deeper into convolutional neural networks for camera relocalization," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 5644–5651.
- [11] B. Wang, C. Chen, C. X. Lu, P. Zhao, N. Trigoni, and A. Markham, "Atloc: Attention guided camera localization," *arXiv preprint arXiv:1909.03557*, 2019.
- [12] Y. Shavit, R. Ferens, and Y. Keller, "Learning multi-scene absolute pose regression with transformers," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 2733–2742.
- [13] T. Naseer and W. Burgard, "Deep regression for monocular camera-based 6-dof global localization in outdoor environments," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 1525–1530.
- [14] A. Kendall and R. Cipolla, "Geometric loss functions for camera pose regression with deep learning," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 5974–5983.
- [15] S. Chen, Z. Wang, and V. Prisacariu, "Direct-posenet: absolute pose regression with photometric consistency," in *2021 International Conference on 3D Vision (3DV)*. IEEE, 2021, pp. 1175–1185.
- [16] S. Chen, X. Li, Z. Wang, and V. A. Prisacariu, "Dfnet: Enhance absolute pose regression with direct feature matching," in *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part X*. Springer, 2022, pp. 1–17.
- [17] A. Kendall and R. Cipolla, "Modelling uncertainty in deep learning for camera relocalization," in *2016 IEEE international conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 4762–4769.
- [18] Z. Huang, Y. Xu, J. Shi, X. Zhou, H. Bao, and G. Zhang, "Prior guided dropout for robust visual localization in dynamic environments," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 2791–2800.
- [19] A. Moreau, N. Piasco, D. Tsishkou, B. Stanculescu, and A. de La Fortelle, "Coordinet: uncertainty-aware pose regressor for reliable vehicle localization," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2022, pp. 2229–2238.
- [20] H. Deng, M. Bui, N. Navab, L. Guibas, S. Ilic, and T. Birdal, "Deep bingham networks: Dealing with uncertainty and ambiguity in pose estimation," *International Journal of Computer Vision*, vol. 130, no. 7, pp. 1627–1654, 2022.
- [21] M. Bui, T. Birdal, H. Deng, S. Albarqouni, L. Guibas, S. Ilic, and N. Navab, "6d camera relocalization in ambiguous scenes via continuous multimodal inference," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*. Springer, 2020, pp. 139–157.
- [22] F. Zangeneh, L. Bruns, A. Dekel, A. Pieropan, and P. Jensfelt, "A probabilistic framework for visual localization in ambiguous scenes," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 3969–3975.
- [23] S. Chen, Y. Bhalgat, X. Li, J. Bian, K. Li, Z. Wang, and V. A. Prisacariu, "Refinement for absolute pose regression with neural feature synthesis," *arXiv preprint arXiv:2303.10087*, 2023.
- [24] H. Blanton, C. Greenwell, S. Workman, and N. Jacobs, "Extending absolute pose regression to multiple scenes," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 38–39.
- [25] S. Brahmabhatt, J. Gu, K. Kim, J. Hays, and J. Kautz, "Geometry-aware learning of maps for camera localization," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2616–2625.
- [26] R. Clark, S. Wang, A. Markham, N. Trigoni, and H. Wen, "Vidloc: A deep spatio-temporal model for 6-dof video-clip relocalization," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 6856–6864.
- [27] N. Radwan, A. Valada, and W. Burgard, "Vlocnet++: Deep multitask learning for semantic visual localization and odometry," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 4407–4414, 2018.
- [28] A. Valada, N. Radwan, and W. Burgard, "Deep auxiliary learning for visual localization and odometry," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 6939–6946.
- [29] A. Moreau, N. Piasco, D. Tsishkou, B. Stanculescu, and A. de La Fortelle, "Lens: Localization enhanced by nerf synthesis," in *Conference on Robot Learning*. PMLR, 2022, pp. 1347–1356.
- [30] T. Ng, A. Lopez-Rodriguez, V. Balntas, and K. Mikolajczyk, "Re-assessing the limitations of cnn methods for camera pose regression," *arXiv preprint arXiv:2108.07260*, 2021.
- [31] B. Glocker, S. Izadi, J. Shotton, and A. Criminisi, "Real-time rgb-d camera relocalization," in *2013 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 2013, pp. 173–179.
- [32] J. Shotton, B. Glocker, C. Zach, S. Izadi, A. Criminisi, and A. Fitzgibbon, "Scene coordinate regression forests for camera relocalization in rgb-d images," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2013, pp. 2930–2937.
- [33] E. Brachmann, M. Humenberger, C. Rother, and T. Sattler, "On the limits of pseudo ground truth in visual camera re-localisation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 6218–6228.
- [34] M. Tan and Q. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *International conference on machine learning*. PMLR, 2019, pp. 6105–6114.
- [35] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [36] T. Sattler, W. Maddern, C. Toft, A. Torii, L. Hammarstrand, E. Stenborg, D. Safari, M. Okutomi, M. Pollefeys, J. Sivic, *et al.*, "Benchmarking 6dof outdoor visual localization in changing conditions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 8601–8610.